AD-A159 914 MICROPROCESSOR LOCAL COMPUTER NETWORK EVALUATION(U) 1/1
GEORGIA INST OF TECH ATLANTA P H ENSLOW ET AL.
15 APR 85 AIRMICS-ISEC-ASBG-84-2 DAAK70-79-D-0087

UNCLASSIFIED F/G 9/2 NL

END

FILMED

DTIC

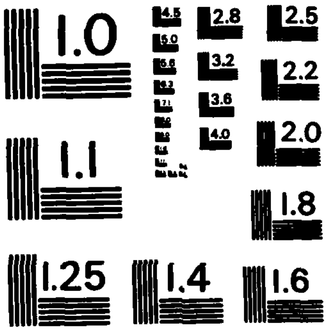

MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS-1963-A

2

AIRMICS

US ARMY INSTITUTE FOR RESEARCH
IN MANAGEMENT INFORMATION AND COMPUTER SCIENCES

USAISC

USAISEC

INFORMATION FOR DECISION

MICROPROCESSOR LOCAL COMPUTER NETWORK EVALUATION

This do[illegible] ed
for pub[illegible]
distribution is unl[illegible]

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS
BEFORE COMPLETING FORM

1. REPORT NUMBER
ISEC-ASBG-84-2

2. GOVT ACCESSION NO.
AD-A159 917

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)
Microprocessor Local Computer Network Evaluation

5. TYPE OF REPORT & PERIOD COVERED
Final Report
1 Jan 83 - 29 Jun 84

6. PERFORMING ORG. REPORT NUMBER
G36-601

7. AUTHOR(s)
Dr. Philip H. Enslow, Jr.
Dr. Joseph L. Hamond
Dr. Jay H. Schlag

8. CONTRACT OR GRANT NUMBER(s)
DAAK70-79-D0087

9. PERFORMING ORGANIZATION NAME AND ADDRESS
Schools of Electrical Eng. Info.& Computer Science
Georgia Institute of Technology
Atlanta, Georgia 30332

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
62725A, SX 762725DY10
04, 008

11. CONTROLLING OFFICE NAME AND ADDRESS
U.S. Army Institute for Research in Management
Information and Computer Sciences
115 O'Keefe Building, GIT
Atlanta, Georgia 30332

12. REPORT DATE
June 1984

13. NUMBER OF PAGES
64

14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)

15. SECURITY CLASS. (of this report)
Unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE None

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

SAME

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Local Area Network (LAN) Emulation, Microprocessor Network, Network Emulation, Network

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes a project that uses a microprocessor local computer network (MLCN) to emulate local area networks which use random access protocols. Specifically, the project is concerned with verifying the ability of the MLCN to emulate local networks which use CSMA/CD (carrier sense multiple access with collision detection). Additional keywords: computer simulation; measurement; packet communications; ETHERNET.

FINAL TECHNICAL REPORT

MICROPROCESSOR LOCAL COMPUTER NETWORK EVALUATION

Philip H. Enslow Jr.
Joseph L. Hammond
Jay H. Schlag

Submitted to

United States Army Institute for Research in
Management Information and Computer Science

Under Contract DAAK70-79-D-0087
GIT Project Numbers E21-610 and G36-601

June 1984

Georgia Institute of Technology
Atlanta, Georgia 30332

TABLE OF CONTENTS

FOREWARD

The work covered by this final report was performed in the Schools of Electrical Engineering and Information and Computer Science. Principal Investigators for the project were Drs. P.H. Enslow, J.L. Hammond, and J.H. Schlag. Contributions to the work were made by H.I. Bibb, D.K. Lam, R. Hutchins, and Y. Tabiri.

Mr. Dale N. Murray was Project Officer for AIRMICS and also a contributor to the work.

I. OBJECTIVES OF THE PROJECT

A microprocessor local computer network (MLCN) test-bed facility was developed at the Georgia Institute of Technology under an earlier project [O'Re83]. One of the features of this MLCN test-bed is the facility for emulating local area networks which use random access protocols. This project is concerned with this application of the MLCN, and, more specifically, with verifying the ability of the MLCN to emulate local networks which use CSMA/CD (carrier sense multiple access with collision detection).

The ability of the MLCN facility to emulate random access protocols accurately has been assessed using simulation data and data from all available published literature. However, there exists only a small quantity of measured data from operating networks. Thus, a thorough evolution of the MLCN facility with respect to actual operating networks has not been possible.

This project was initiated to address the problem of evaluating the MLCN with respect to an operating network. The project had two general goals namely: 1) to obtain measured data and develope actual operating characteristics from an in-service random access type local network, and 2) to use this data to evaluate the ability of the MLCN facility to emulate the behavior of such a network.

The choice of operating network for the project was limited to what was available, namely an Ungermann-Bass Net/One network. Unfortunately, the structure and characteristics of Net/One are very poorly documented for two reasons. First, Ungermann-Bass regards many features of the network as company confidential and secondly, measurement of operating characteristics is a non-routine task and certainly a nontrivial one.

Achieving the first of the project goals required a number of tasks which included: design of measurement software and hardware, design of experiments, obtaining appropriate data, and constructing models for the

Net/One for several conditions of operation.

Given measured characteristics of Net/One and a model for its operation under varying conditions, the second part of the project is to emulate the structure of Net/One on the MLCN facility and measure essentially the same data as was measured on the operating network. A comparison of these measured results would serve to evaluate the accuracy of the emulator.

It was clear at the beginning of the project that the approach envisioned entailed certain limitations. The most important of these are: limitations imposed because of available equipment (e.g. the Net/One installation has something less than one dozen stations), the fact that measurement and modeling of Net/One had to take place before emulation could be carried out, and the fact that the complexity of the Net/One architecture was not known at the inception of the project. As is discussed in the conclusions to this report, all of these factors affected the breadth of the study.

2. SUMMARY OF TYPICAL RANDOM ACCESS NETWORK ARCHITECTURES

To discuss the project in the proper perspective it is desirable to provide some background on typical architectures for random access networks. Two organizations have been involved in developing standards for networks, namely the International Organization for Standardization (ISO) and the Institute of Electrical and Electronic Engineers Project 802 Standards Committee (IEEE 802).

To minimize design complexity, modern networks tend to use a layered architecture with each layer being a logical entity which performs certain functions. The services performed by the highest layer are for the network users. Lower layers each provide services for the next higher layer, shielding it from the details of how the services are provided. This hierarchical structure can be used to provide all of the intermediate tasks and functions required to carry out desired user-to-user interaction over the network.

When users located at different network nodes communicate, the corresponding layers at each node also communicate. A well-defined set of rules called a "protocol" is necessary for each level to carry out its conversation in an orderly structured manner. These protocols establish what can be regarded as logical communication paths between communicating entities, which are called "peer processes".

Protocols are organized in a hierarchial fashion to correspond to the network layers. All but the lowest layer protocols control conversations across a single layer over a logical, or virtual, communication path. The lowest layer protocol controls data flow over the physical connecting channel.

The number of layers, the name of each layer, and the function of each layer differs from network to network. The two extreme layers are usually the application, (or user-user), and physical layers. The application layer, as its name implies, deals with the specific application of the user. It is always the highest level. The physical layer, on the other hand, is always the bottom layer, as it deals with the transmission of physical signals between the two nodes. The physical layer includes the physical trnsmission channel and provides an electrical connection between the communicating processes.

Two types of network services typically available are termed "datagram" and "virtual circuit" in the ISO literature. Datagram service is the more basic of the two just named. In providing such a service, the network and lower layers couple the sender and receiver by transferring packets in a completely independent manner over the virtual connection. Datagram service can be likened to the postal service in that packets are delivered to a destination address, which must be specified in the packet, but such delivery constitutes all that is required. No error control or end-to-end acknowledgements are provided, and datagrams are passed to the user-host in the order in which they arrive.

As can be seen from the above description, datagram service is fairly primitive. Virtual-circuit service is at the other extreme, providing all of the functions noted above as being absent from the datagram service. Thus the virtual-circuit service not only delivers each packet to its destination but also provides on an end-to-end basis proper sequencing of packets, error control, flow control, and acknowledgements.

Figure 1 shows the levels identified in the ISO model and those specified by the IEEE 802 committee. Note that the seven layers in the ISO model can be divided into end-to-end protocols and network access protocols.

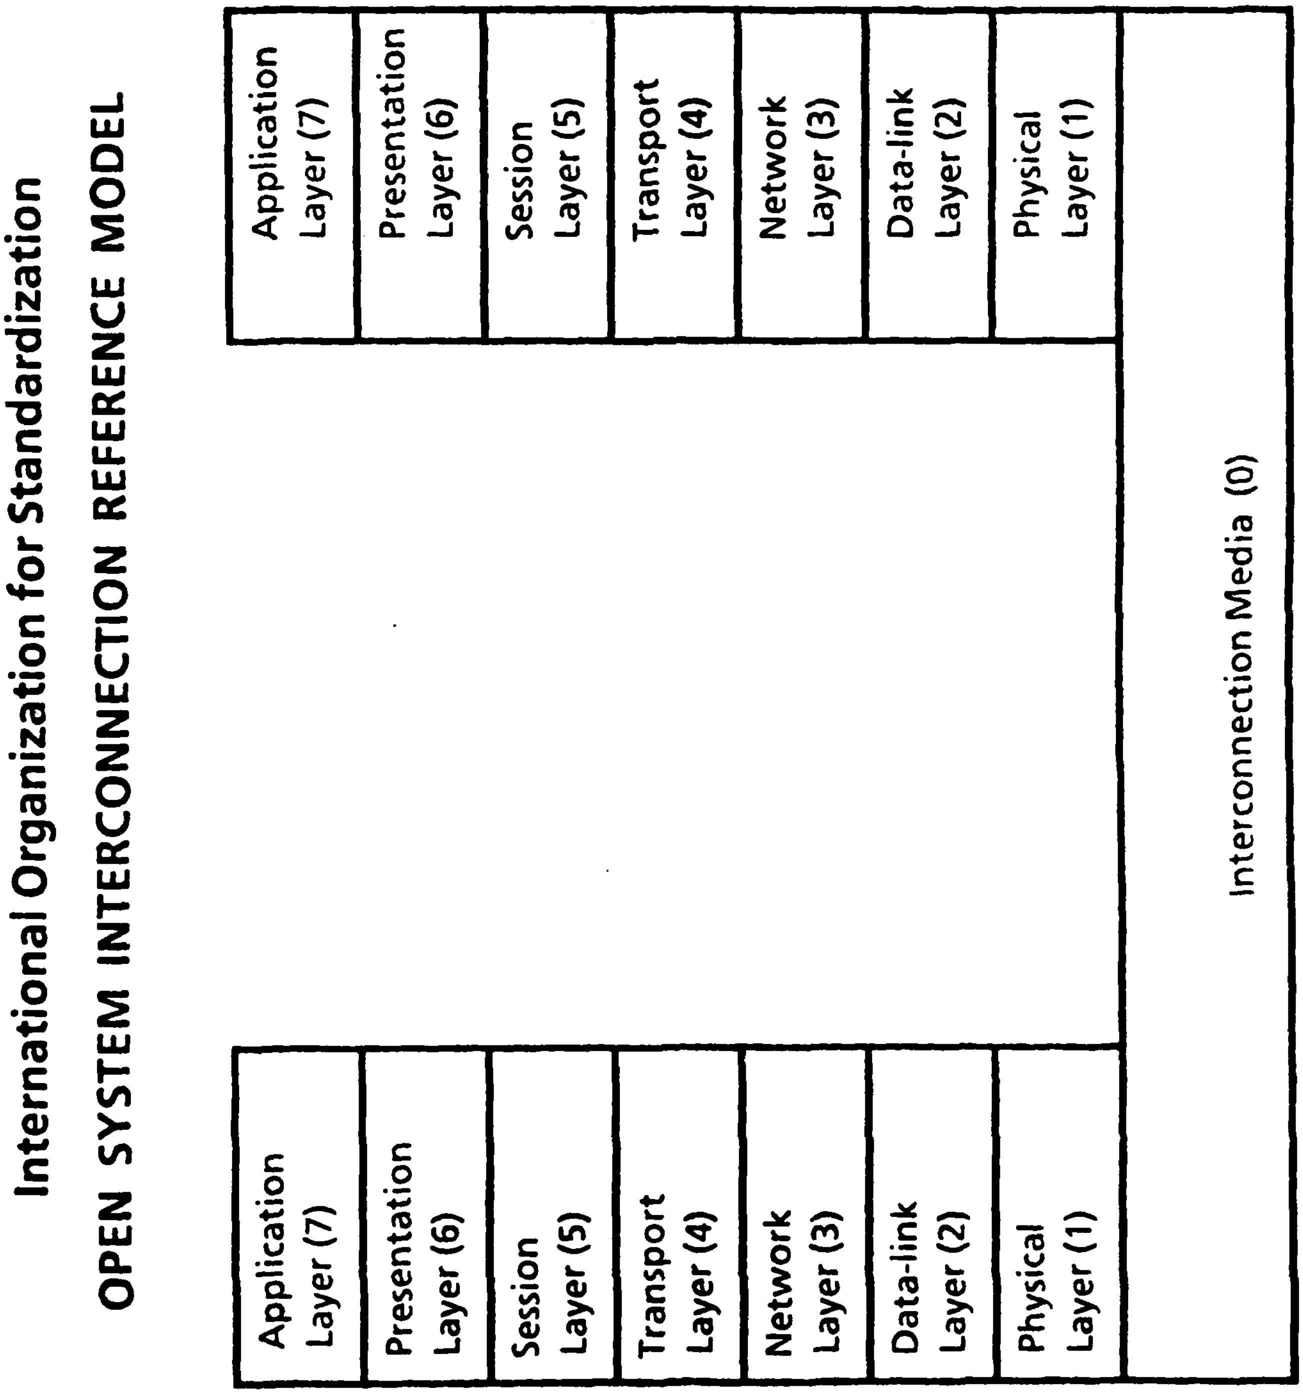

Figure 1

The ISO Model and the IEEE Local Area Network Reference Models Compared

In reviewing the layered structure it should be noted that the IEEE 802 standard layers - physical, media access control (MAC) and logical link control (LLC) - encompass esentially all that is needed for network access. For the ISO model, the group of layers for network access provide all of the necessary features and functions required to provide basic datagram or virtual circuit services, although actual delivery of these services is a function of the transport level in the ISO model. Most additional services, provided by levels higher than the network access levels, are no longer in the area of basic communication but fall into application specific tasks or services which are not of interest to this project.

3. PRELIMINARY DESIGN OF EXPERIMENTS

The basic operation of a local area network such as the Net/One can be viewed in terms of the datagram or virtual circuit services that it provides. As discussed in Section 2 of the report, both types of service involve the transfer of packets between source and destination nodes. The network performance in transferring packets can, in turn, be characterized by relating throughput and average delay to to offered traffic.

Offered traffic is defined as the average input traffic rate, in packets per second, to each node. Total offered traffic is the sum of the average input traffic to all nodes on the network. Throughput as used in this report, is a normalized measure of average output traffic either from some specific node or from the complete network. It is defined as the average output traffic, in packets per second, from some portion of the network divided by the maximum packet rate of the output channel. The packet rate of the output channel is the bit rate or clock rate of the channel divided by the average packet length.

As a final definition, the average delay for a packet is the delay in seconds, averaged over a reasonable length of time, for a packet to traverse a given portion of the network. The average delay normally consists of waiting time in a buffer, propagation delay and the packet transmission time required after processing of the packet begins. Average delay can be measured over any portion of the network. End-to-end delay and round-trip delay are two types of delay often cited.

The major thrust of the subject project can be stated in terms of the variables defined above as follows: 1) For Net/One operating under a set of specified conditions, measure the throughput and average delay for a sequence of offered traffic values, 2) Use this and other data to establish the structure of Net/One, 3) Emulate the structure of Net/One on the MLCN facility, and make measurements of throughput and average delay which correspond to this mode for Net/one, 4) Compare the results. One of the tasks of the project was to design the experiments by choosing the conditions under which throughput and delay are measured.

The general structure of the Net/One is shown in Figure 2. The end-to-end average packet delay between two user devices can be broken down into a delay for each of the components of the path identified as network processor, transceiver interface, etc. An alternative breakdown of delay into components would use the layers discussed in Section 2. Classified in this way, using the IEEE 802 notation results in the following components: delay due to the physical transmission medium (cable), media access delay, logical link control delay and delay due to highr level protocols. The first two layer delays are asociated with the cable, the transceiver and the transceiver interface in Figure 2, while the others are associated with the network processor.

The MLCN facility was designed to emulate directly the physical transmission medium and the media access protocol.

At the inception of the project it was planned to duplicate the significant aspects of the Net/One network processor, the operation of which is primarily software, with software running on two MLCN computers which normally serve as the physical stations for the MLCN facility. (See References [1] for a complete description of the MLCN facility.) The approach would result in an emulation for station-to-station communication in Net/One as shown in Figure 3.

A consideration of the overall emulation problem or the related problem of characterizing Net/One led to a two-step approach. In a first step the delays and characteristics of the physical and media access layers would be measured and emulated. This would then be followed by the second step of working with complete end-to-end communication between two users coupled to different stations. The remainder of the report is structured, to a large extent around these two steps in the procedure.

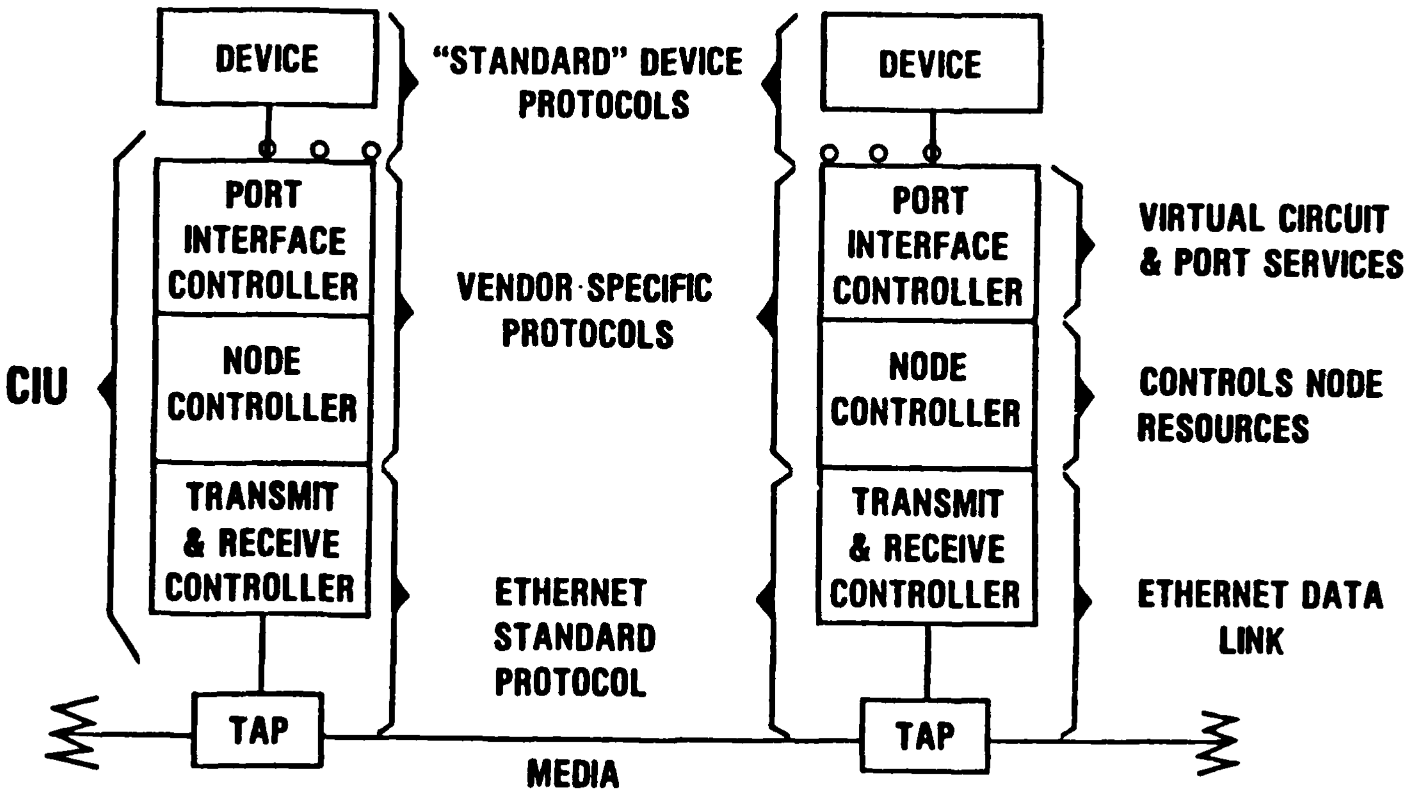
DEVICE
"STANDARD" DEVICE PROTOCOLS
DEVICE
CIU
PORT INTERFACE CONTROLLER
NODE CONTROLLER
TRANSMIT & RECEIVE CONTROLLER
VENDOR·SPECIFIC PROTOCOLS
PORT INTERFACE CONTROLLER
NODE CONTROLLER
TRANSMIT & RECEIVE CONTROLLER
VIRTUAL CIRCUIT & PORT SERVICES
CONTROLS NODE RESOURCES
ETHERNET STANDARD PROTOCOL
ETHERNET DATA LINK
TAP
TAP
MEDIA

Figure 2. General Logical Structure of the Net/One Local Area Network.

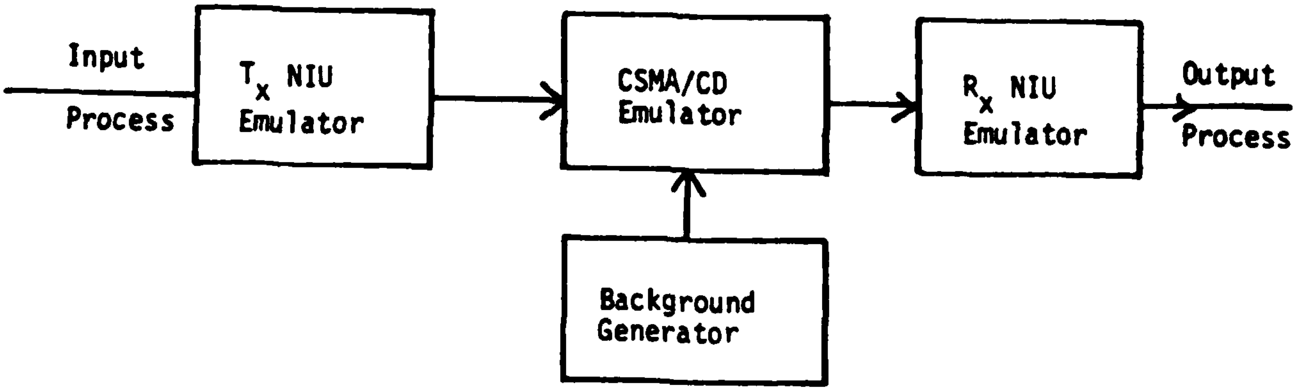

Figure 3. Emulation of Station-to -Station Communication for Net/One.

4. NET/ONE PARAMETERS TO BE MEASURED

In order to fully characterize the operation of the LAN network under study, it was necessary to accurately determine the value of a number of parameters describing the operation and performance of the system. The specific parameter values required were:

● Size of the data packets --- the size of the packets placed on the cable and the number of device data characters per packet (not the total size of the ETHERNET packet data field). In our LAN these range from 1 to 64 data characters, with flow control invoked at 42.

● Delays/time intervals

●● End-to-end character delivery delays

●●● Packet building delays

●●● Protocol delays

●●● Cable delays

●●●● Access delays

●●●● Transit delays

It was necessary to obtain values for these parameters under various loading conditions with respect to both device speeds and number of active devices being serviced as well as the cable loading. All of these parameters are random variables whose underlying distributions were not known and could not be assumed. This made it necessary to capture and record a large number of observed values and analyze the resulting histograms depicting their distribution.

4.1 General

The Ungermann-Bass Net/One is a 10 Mbs network utilizing the Ethernet specification at the physical and data link layers. The primary user interface is virtual circuit service provided by Z80 processors operating in a multiple processor configuration. Each node of the network is housed in a single box containing up to four processor boards, a transmitter board and a receiver board all operating as one unit called the Network Interface Unit (NIU).

Each NIU consists of one Network Processor Board (NPB) and up to three Application Processor Boards (APBs). The APBs are all optional, but all the NIUs used in these experiments had three. The user ports are on the four processor boards with six ports on a board. The NIU is attached to a radio-frequency transceiver. The transceiver is connected directly to the coaxial cable transmission medium. Figure 4 illustrates the Net/One NIU organization.

The primary user services provided are interfaces to asynchronous, start-stop, character-mode DTE's, terminals, computers, or other perpherals. ETHERNET [DIX 80] is utilized to provide the media access control mechanism and lower level protocols. Since the ETHERNET system is fundamentally a packet interconnection system, the individual characters transmitted by the attached devices must be assembled into packets before transmission (and disassembled prior to delivery). In addition, the ETHERNET specification provides only "unreliable datagram service," while the interconnection of two devices require a reliable connection service. Therefore, it is necessary for the NIU's to also include a Virtual Circuit Protocol.

4.2 The Packet Building and Buffer Management Processes

The attached devices send individual characters utilizing the asynchronous, start-stop protocol. As each character is received, it is processed by the serial interface chip and placed in a port-buffer in the Interface Controller memory. Associated with each device interface port is a pool of buffers utilized to assemble characters into packets for transmission over the media. The packet building and buffer management processes are the most important activities effecting the performance of the LAN system under study. Figure 5 illustrates the movement of data and the transfer of buffers from one activity to another. Figure 6 illustrates the timing relationships between the various activities that take place and gives an overall view of the network operation.

As in almost every other similar system, buffer management is a critical issue. In this particular system, dedicated buffers are assigned to each port. The operational scheme is depicted in Figure 5. There are three buffers assigned to each port, and all ports function independently (up to a point) and in parallel. A BUFFER-SERVICE-INTERRUPT occurs upon the expiration of a timer in the Processor Board (16.4 milliseconds). When this interrupt occurs, the processor checks to determine if it is possible to close off the current buffer and move it to the SEND-QUEUE for transmission to the other end of the virtual circuit. A port buffer is closed and moved to the SEND-QUEUE if either of the following conditions apply:

- There is another empty buffer immediately available to continue to receive input characters, <u>or</u>
- The current buffer is completely filled.

After the buffer has been closed, the virtual circuit header and the ETHERNET header are added to the data character(s). It may be necessary to add up to eight filler characters to meet ETHERNET specifications for a minimum-size packet.

Figure 6 is a time-sequenced chart of what happens to an individual packet as it moves through the LAN from one NIU to another. One cannot help but notice a similarity between progression of data through the NIU and movement of data through a classical store-and-forward type of network. It should not be surprising then that most of the same design and performance issues are present. These are flow control, congestion, addressing, and throughput problems. Even routing issues are present as the virtual circuit addressing is processed.

The completed ETHERNET packet is transferred to the SEND-QUEUE for that Interface Controller by moving a pointer, as depicted in Figures 5 & 6. The

SEND-QUEUE is a FIFO (first-in, first-out) service order queue. There are separate SEND-QUEUEs for each port controller board (i.e. APB or NPB) in the system. There can be up to four boards in this product. Each controller board may have up to six ports. The node controller continually polls all of the Processor Boards to determine if there is a packet ready to be transmitted. There is a classical random queuing delay at this point that has several distinct parts. The buffer must wait until it reaches the head of the queue. This delay is a function of the number of buffers ahead of it, the service or transfer time for a buffer, and the wait for the next poll by the Node Controller to service the SEND-QUEUE of a particular controller board. From one of the SEND-QUEUEs, on a port board, the packet is copied to the TRANSMIT-BUFFER on the node controller board.

In this system the copy is by DMA transfer. The buffer cannot be released from the SEND-QUEUE for reuse at this time, due to the guaranteed delivery or Virtual Circuit requirement. It must be held here until either an acknowledgement is received or a timeout occurs triggering a retransmission of this packet.

The packet in the TRANSMIT-BUFFER now experiences another random delay. When the packet reaches the head of the TRANSMIT-BUFFER queue, the random service delay has several components of significance. Here is where the access control mechanism comes into view. If it is contention based (ETHERNET, as our case) the cable access, collision resolution (random, truncated binary-exponential delay), and retransmission delays must be counted. The final component of this time delay is the transmit time which is a linear delay based on 10 MBS cable speed and ETHERNET packet specification. The range is from 67.2 microsec minimum packet size up to 1.23 milliseconds for a maximum 1500 octet packet.

After the ETHERNET packet has been successfully received, the complete packet is placed in the RECEIVE-BUFFER in the receiver. The ETHERNET header must be retained for the higher level acknowledgement, so the size is not reduced. The node controller processes the RECEIVE-QUEUE in FIFO order. Another classical queueing delay is experienced. The virtual circuit header is checked to determine where to move the buffer. This processing time is still random, but well behaved as compared with some of the others. It is a function of the total activity of the Cable Interface Unit and the competition for CPU cycles. Once the location is determined, the packet is copied into a RECEIVE-QUEUE on a port controller board. Net/One also uses DMA for this transfer. Again random queueing and service delays are experienced. There is a wait for the DMA to be set-up, and then the copy time is a function of the packet length and the CPU-cycle competition.

In the RECEIVE-QUEUE, the virtual circuit header is examined to determine which output port is to receive the data, and where to return the acknowledgement. An ACK-packet is constructed and set up for transmission back to the sending controller board. Meanwhile, the buffer is moved by a pointer change from the RECEIVE-QUEUE to a particular port controller. It is serially transmitted to the destination device at a data rate determined by that device and subject to its flow control. The delays are proportional to the bit rate and the length of the raw data with all headers removed. The acknowledgement packet traverses the system in the same manner just described, back to the originating port controller board. Figure 6 shows the additional time for processing the return ack so that the buffer can be released for reuse.

4.3 Data structures

A large number of data structures are required to handle all 24 virtual circuits on the four processor boards in an NIU. Each processor board has one transmit control block (XCB) and an array of thirteen data type descriptor control blocks (DTDs). Each XCB consists of a control word used to determine if the XCB is in use and valid, an address field used as the source of the current packet and an address for the start of the next packet. A pointer points to one of the buffers on the send queue, which is a linked list of packets ready to be sent. The DTDs contain information similar to that in the XCB in addition to the linked list of fill and receive buffers into which the NPB can perform direct memory access (DMA) of packets addressed to the associated DTD. Fill buffers are used by the NPB as the destination of DMA transfers from the receiver board. Receive buffers are seen by the receiving process as the next to empty. All data structures are stored in shared memory to permit the NPB to access them. A process, in order to be able to have access to the transmission medium, must, on initialization, request a data type from the network manager. It can request a specific data type for a special purpose or can request for the next available data type. Thus, data types are associated with processses and are addressed by the processes to which they are assigned.

4.4 NIU Packet Transmission Process

All data transfer out of the NIU is handled by the transmitter process on the transmitter board. The transmitter process consists of a 4K byte FIFO buffer, direct memory access (DMA) logic to enable the transmitter to obtain data directly from another processor's memory, and the logic to detect collisions. Collisions are detected by the transmitter monitoring all data transmissions and doing a bit by bit comparison between the transmitted and received data. The interface between the NPB and the transmitter process is provided by the Z80 I/O ports which addresses command registers on the transmitter board. The NPB scans the XCBs of the processor boards. If it finds an XCB which is ready and valid, it sets up a DMA transfer directly to the FIFO of the transmitter board by transfering to the transmitter register the source address, the byte count and the DMA enable command. DMA then begins, and is controlled by the transmitter process until the byte count reaches zero. At the zero byte count, the transmitter issues an interrupt to the interface process running on the NPB. The NPB then enables the transmitter to transmit the packet as the transmission medium permits. As soon as the packet has been successfully transmitted, the transmitter process notifies the NPB. The XCB of the transmitted packet continues to be marked as in use until an acknowledgement to the packet has been received. Figure 7 illustrates the NIU transmission process.

4.5 NIU Packet Receiving Process

Receiving a packet is done by the receiver process which runs on the receiver board. The board contains a 4K bytes of FIFO buffer, DMA logic, command registers, and a 32-bit cyclic redundancy check (CRC) logic. The receiver automatically checks the Ethernet header on all packets to determine if the packet is addressed to its node. If it is not, the receiver ignores the packet and continues to sense the carrier. Otherwise, the receiver begins at the current location in its FIFO and loads the data into it. When the end of the packet is reached, the FIFO pointer is updated to reflect the current location. The receiver next interrupts the NPB, which determines the processor board to which the packet is addresssed. If the packet is good, the NPB checks the DTD of the destination processor board to determine if a buffer is available. If not, the packet is discarded. Otherwise, a DMA transfer from the receiver FIFO is set up to the fill buffer and a control word is set to indicate that tha buffer is in use. The NPB then interrupts the processor board, which polls its DTDs to determine which one has a ready packet. Any ready packet is removed from the fill queue and placed on the receive queue. The APB or NPB process has the responsibility to notify the receiving process of a packet arrival and returning the empty buffer back on the fill queue. However, it is the responsibility of the receiving process to process the data in the packet and flag the receive buffer as free. Figure 8 illustrates the NIU receiving process.

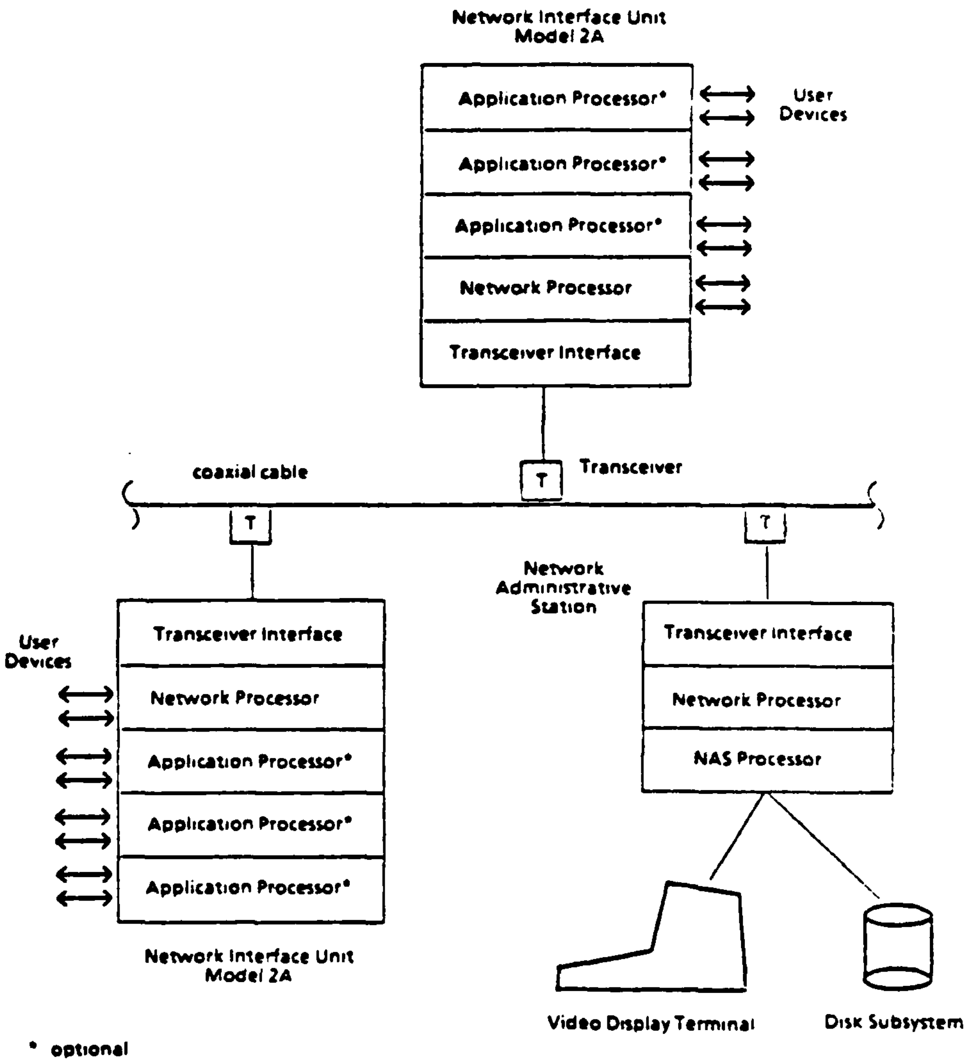

Figure 4. Net/One Hardware Architecture.

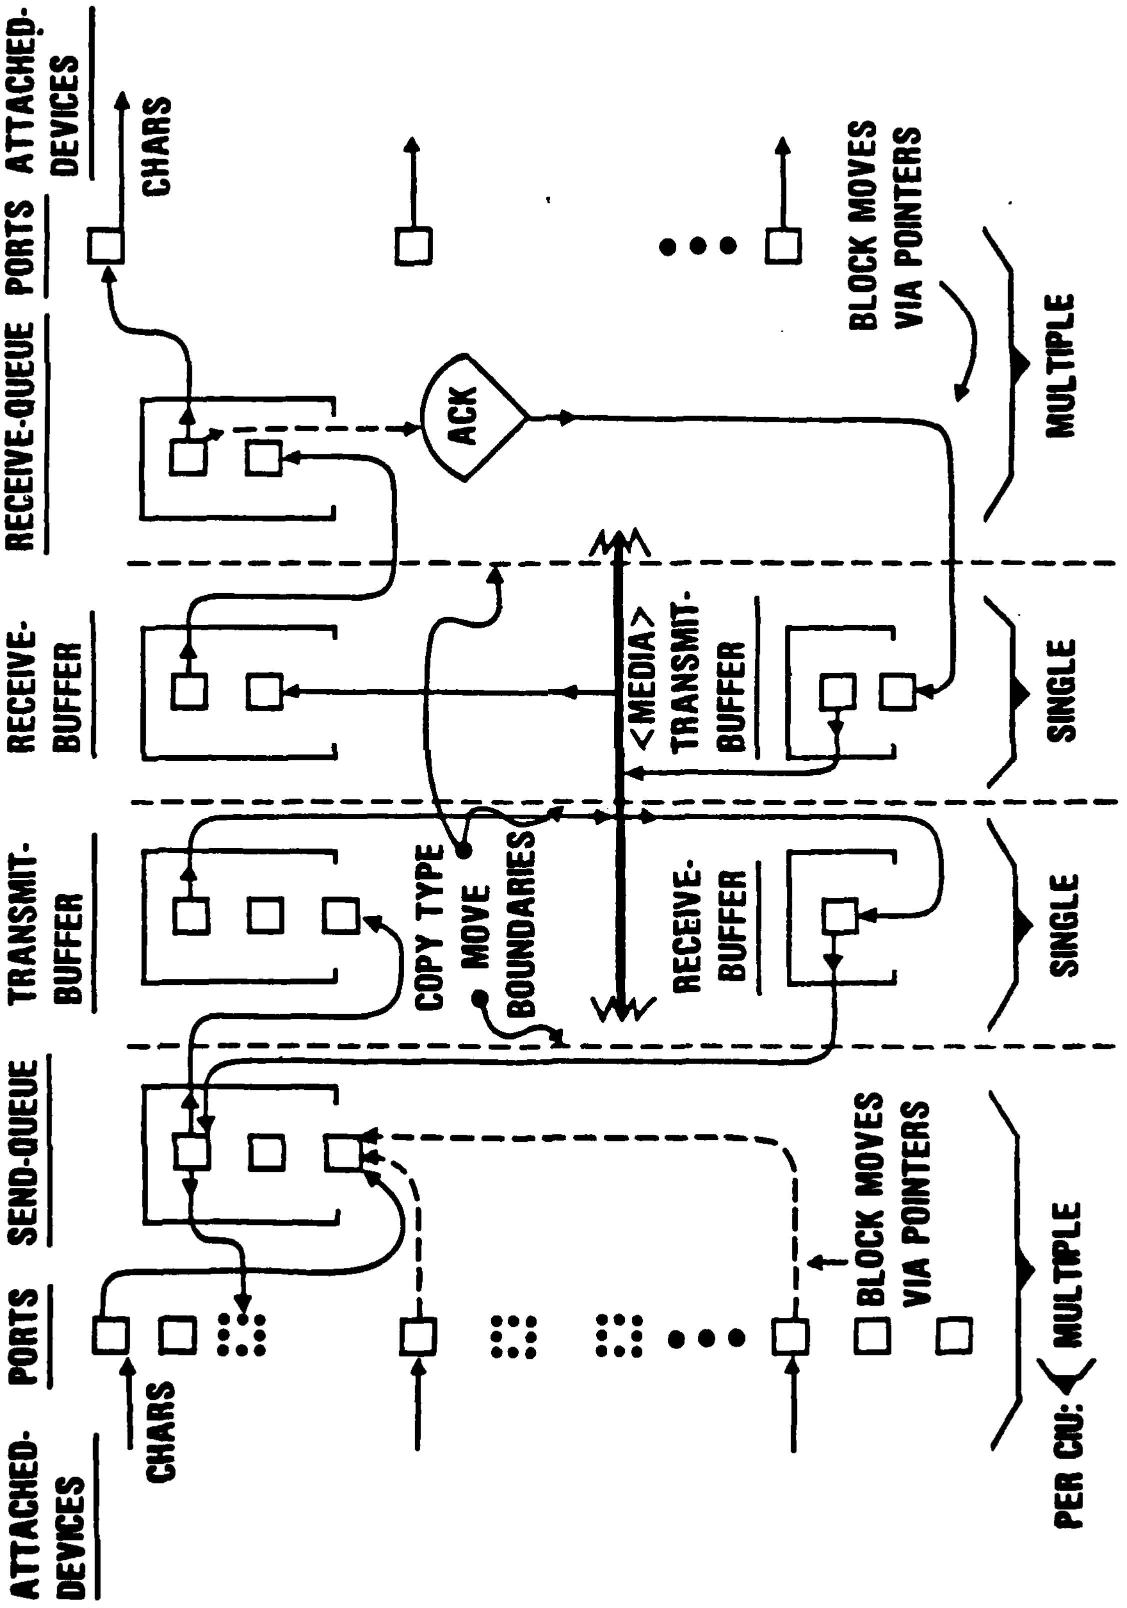

Figure 5. The Buffer Management Process.

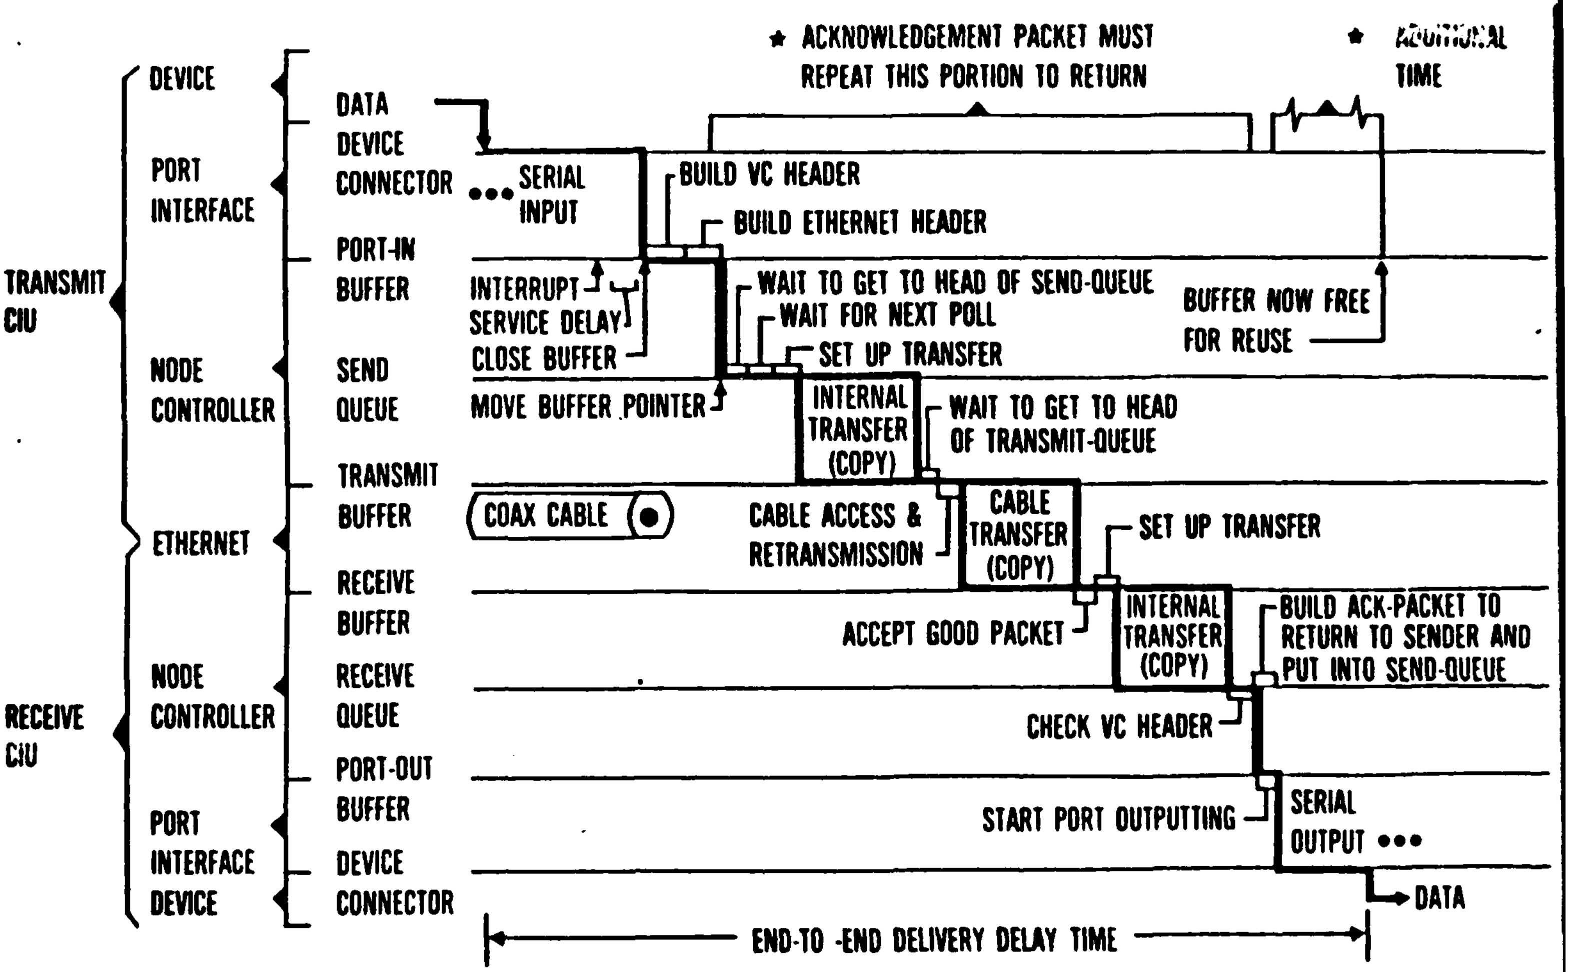

Figure 6. Time Chart: Packet Building and Virtual Circuit Operations.

Figure 7

NIU Packet Transmission Process

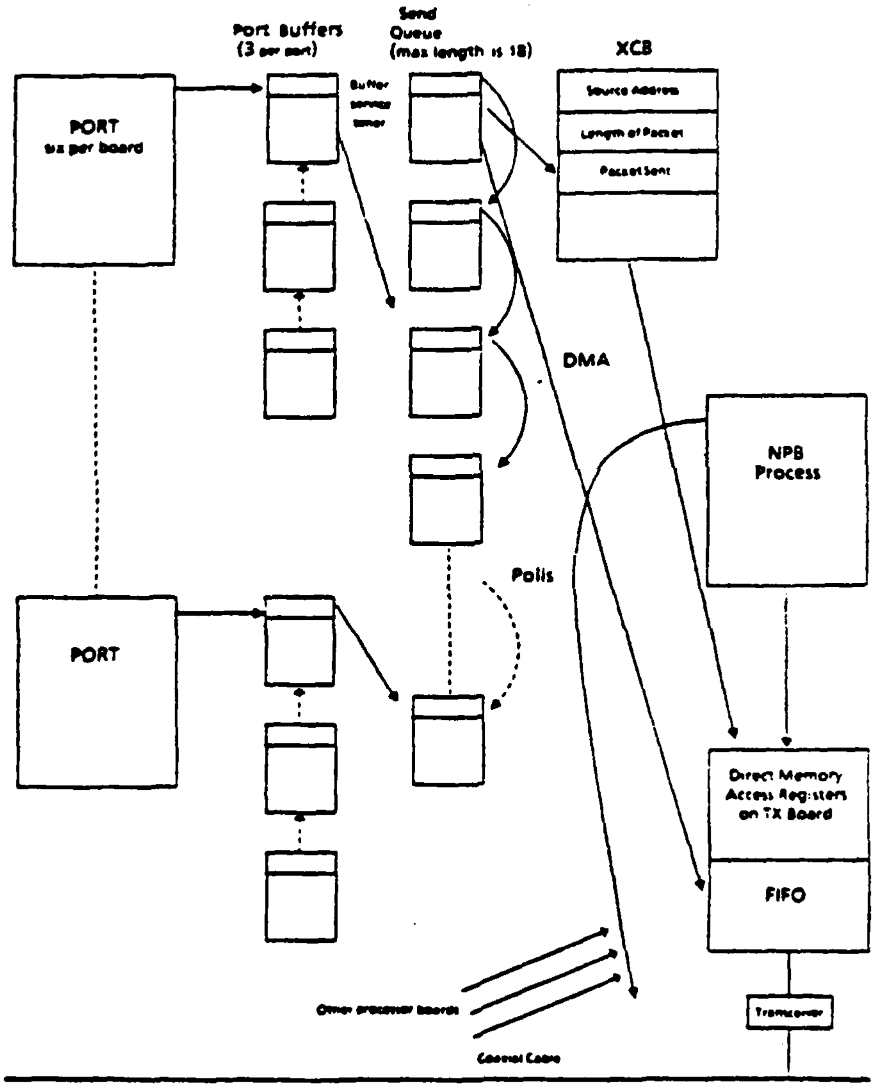

Figure 7. NIU Packet Transmission Process.

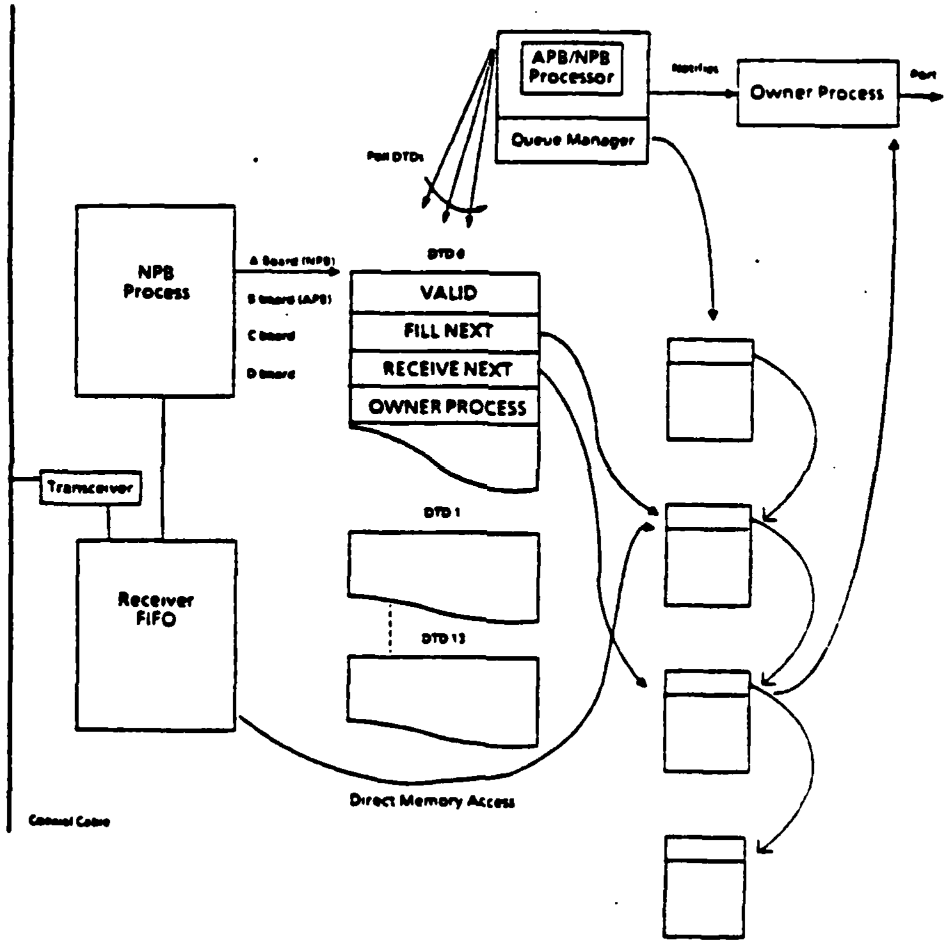

Figure 8. NIU Packet Receiving Process.

5. MEASUREMENT TOOLS AND TECHNIQUES

The experiments were performed in the Computer Laboratory of the School of Information and Computer Science at the Georgia Institute of Technology using the Net/One facility described above.

Initial experiments were designed to measure the performance of the physical and media access layers in the Net/One architecture. These layers satisfy what is often termed the "Ethernet" specification.

To carry out this part of the study all elements of Net/One, as shown in Figure 2, above the transmit and receive controller were deactivated. Traffic was generated by specially designed software and injected into the transmit controller.

Experiments consisted of injecting traffic, with either a constant or exponentially distributed time between arrivals, into the transmit controller on each of the five stations. A monitor station was then used to measure network throughput as a function of input traffic load. The "load" was determined by adjusting the average time between packet arrivals.

Experiments subsequent to the initial one were performed with Net/One in its normal configuration. Tools and techniques for these experiments are now described.

A major concern in our study of the Net/One performance was understanding how the NIUs built Ethernet packets from the input data and determining the size distribution of those packets under various user input and cable loads. A significant group of experiments examined the process of packet building in the NIU port buffer. Since backoffs from a busy cable could have a 'back-pressure' effect, an empty cable was used initially. This isolated the operation of port buffer management so that its behavior was affected only by the ability of the board processor to handle incoming traffic.

5.1 Traffic Generator

Traffic was generated with a character generator adjustable to traffic rates of 109, 291 and 436 characters per second. The traffic stream could be directed into any number of ports simultaneously. Since each board has its own dedicated processor, the use of the same data stream for simulating user input traffic should not affect the timing for the buffers. Measurement of packet size were made at another NIU (the monitor), after the lower level protocol wrapper (i.e., Ethernet and Net/One virtual circuit header), had been stripped off. As each packet arrived, the size field in the Net/One header was examined and a counter in an array indexed by the packet size was incremented. Each experiment lasted five minutes, long enough for the network to stabilize. Packet sizes of 37 bytes were ignored; they were acknowledgement frames. All packets shorter than 37 bytes were packet fragments caused by collisions and were also ignored.

5.2 Packet Length Measurements

Although it might have been possible to obtain data on the length of the packets generated by an active Cable Interface Unit through direct measurement of its operations, the taking of such data would probably have distorted the operation of the Controller and made the figures obtained inaccurate. The method chosen was to utilize a separate Network Interface Unit operating in the permissive mode, i.e., receive all packets. This separate NIU permitted the examination of the virtual circuit header and recording of the values of the length fields in those headers. This data was collected and presented in the form of histograms.

5.3 Character Delivery Delay Measurements

A special measuring device was constructed to measure the character delivery delays. This measurement device has the capability to send characters to the Network Interface Unit at selectable constant rates. At the same time it can determine the delivery delay from NIU to NIU within a half millisecond using a timestamp technique. The device collects the data and presents a histogram depicting the distribution of the delivery delays. Utilizing this measurement device, it was possible to run a wide variety of experiments for a length of time sufficient to obtain high statistical confidence in the distributions obtained.

6. MEASURED RESULTS

6.1 Results for "Ethernet Type" Operation

Measured data for the initial experiment, involving only the physical and media access layers of Net/One, is plotted in Figure 9 as throughput versus packet arrival rate at each station. Five stations were used, each with the same arrival note. Operating in the somewhat artificial mode, as described in the last section, it was possible to load the network to a throughput of 0.6 as shown in Figure 9.

6.1 Packet Length

Figure 10 shows the distribution of packet sizes with one active port on the APB and a character arrival rate of 109 character per second. This arrival rate corresponds to about 1.78 characters per time slice of 16.4 milliseconds, specified in the Net/One description as the length of the port buffer cut-off timer. In Figure 10, it can be seen that the mean of the observed packet sizes is about 2.54 data bytes. This indicates that buffers were closed at approximately 18.4 or 27.6 milliseconds time interval. This suggests that there are some secondary delays after the interrupts before the processor can close the buffer and move the packet to the send queue.

Figure 11 shows an exponential distribution of packet sizes. There were four active ports on an APB,each with a character arrival rate of 109 per second. As in Figure 10, the character arrival rate per time interval of 16.4 milliseconds per port was 1.78. This suggest that as the number of active ports increases, the service distribution of the processor becomes more random. About 90% of the packets are between 2 and 22 data bytes. With a mean packet size of about 6 data bytes, buffers are serviced on an average of about 64 milliseconds time interval This time of about 64 milliseconds is about four times the buffer time interval (16.4 ms), suggesting that the processor is becoming heavily loaded servicing the port buffer.

Figure 12 is almost identical to Figure 10. Both figures have similar characteristics. The only difference between the two is that Figure 10 used an APB while Figure 12 used an NPB. The similarities in both figures suggest that at low character rates, there is no difference in the packet size distribution for an APB or an NPB.

In Figure 13 the packet size distribution follows a normal distribution with a mean of about 5 data bytes. This distribution is in contrast to Figure 11. As may be recalled, the only difference between the measurement characteristics of Figures 11 and 13 is the type of processor board used. The packet distribution pattern of Figure 13 suggests that as the number of ports increases, the buffer servicing rate on the NPB becomes faster than that on the APB. Using a mean packet size of 5 data bytes, the average buffer service interval is about 45.9 milliseconds.

Figure 14, which depicts the packet size distribution with one port active at 291 characters per second (cps) on an APB, shows packets distributed equally between sizes of 8 and 9 data bytes. With a mean packet size of 8.5 data bytes, the buffer timer is about 29.2 milliseconds, again almost twice the length given in the Net/One specification. This distribution, that is, all packets centered about two values indicates availability of a free port buffer at all times. There was an avilable buffer because there was only one port active in the entire NIU and the data rate was moderate. As the number of active ports is increased to 4 as shown in Figure 15, buffers were open for longer time, resulting in larger packet sizes. This suggests that as the number of ports increases, the servicing time for the port buffers becomes longer.

The mean packet size in Figure 16 is 9 data bytes. This is almost identical to Figure 14. The average buffer length of the service time corresponding to this distribution is about 31 milliseconds. Again, this repeats that buffer timing is almost double the number given in the Net/One specification. Figure 17 sh the same behavior as Figure 18 and can therefore be explained similarily.

Figure 18 shows one active port on NPB at a character arrival rate of 436 per second. This arrival rate corresponds to 7.15 characters per the 16.4 milliseconds time interval. From the figure, packets are distributed around multiples of 14, that is 14, 28, 42, and 56 data bytes. This packet distribution means that buffers were closing at time intervals of approximately 32 milliseconds. This packet size distribution indicates that the length of a packet greatly influences its delay in the port buffer. Thus, high character arrival rates results in longer packet, which in turn increases buffer servicing timing.

When the different character rates are compared, there is a trend towards a higher minimum packet size. As noted before, at 109 cps there are almost no single byte data packets, but about 90% are of 2 or 3 data bytes. At 291 cps the minimum data packet is 8 or 9 bytes, and at 436 cps (Figure 19) the smallest peak is at approximately 17 data bytes. The size of the data bytes seems to depend on the number of characters arriving during a time interval, but it is almost twice the expected amount based on the 16.4 milliseconds time value.

Table 1. Smallest significant data size observed and data sizes that will be produced assuming 16.4 ms port buffer service interval with one active port.

CPS	Data size per 16.4 ms	Smallest significant packet size observed
109	1.79	2 or 3
291	4.77	8 or 9
436	7.15	17 on the average

The second group of experiments (Figures 20 and 21) examined the packet size distribution for a typical traffic load on the network. The group at one data byte probably results from low speed human input. The peak at 15 data bytes indicates a process that produced constant character rate of 468 assuming 32 milliseconds timer for port buffers. The packet sizes of 70 data bytes were probably produced by a Xerox Star system which was also operating on the same cable during this experiment.

The third group of experiments involved traffic generated by a typist with a background load simulated by a special software packet generator. The expected result is that as the load on the cable increases the send queue becomes longer resulting in long delay in servicing port buffers. This will therefore increase packet sizes. But the results, as shown in Figures 22, 23 and 24 do not indicate any significant change in packet sizes as the load on the cable increased. In view of this, the conclusion was that the cable may be so much faster than the NIU that its load is not a factor in the NIU operation.

6.3 Delivery Delay Measurements

As should be expected, delivery delays showed the same multi-mode distributions as packet lengths. Figure 24 is the distribution of delays for one port active at 1200 bps. Most experiments have shown a single peak around 81 ms, whereas this particular one has only 34% there, with 52% clustered around 61 ms. The rest of the data is scattered with the largest being 3.6% at 66 ms. We are currently trying to analyze the behavior of these delay distributions.

Figure 25 shows the distribution of delays for one port at both 1200 and 4800 bps. For the 4800 bps experiment, there is 59% clustered at 79.5 ms, 40% at 77 ms, and only 1% scattered. These two peaks are separated by approximately one character time and represent the one character uncertainty.

For these experiments we collect data for about 300 seconds for each run. Ahtough the packet length and delay time measurements are not synchronized, they are done approximately at the same time and results are reasonably repeatable. Not every packet is measured at the faster bit rates, but a random sampling is done. For both experiments the number of delay measurements is on the order of 36000 packets.

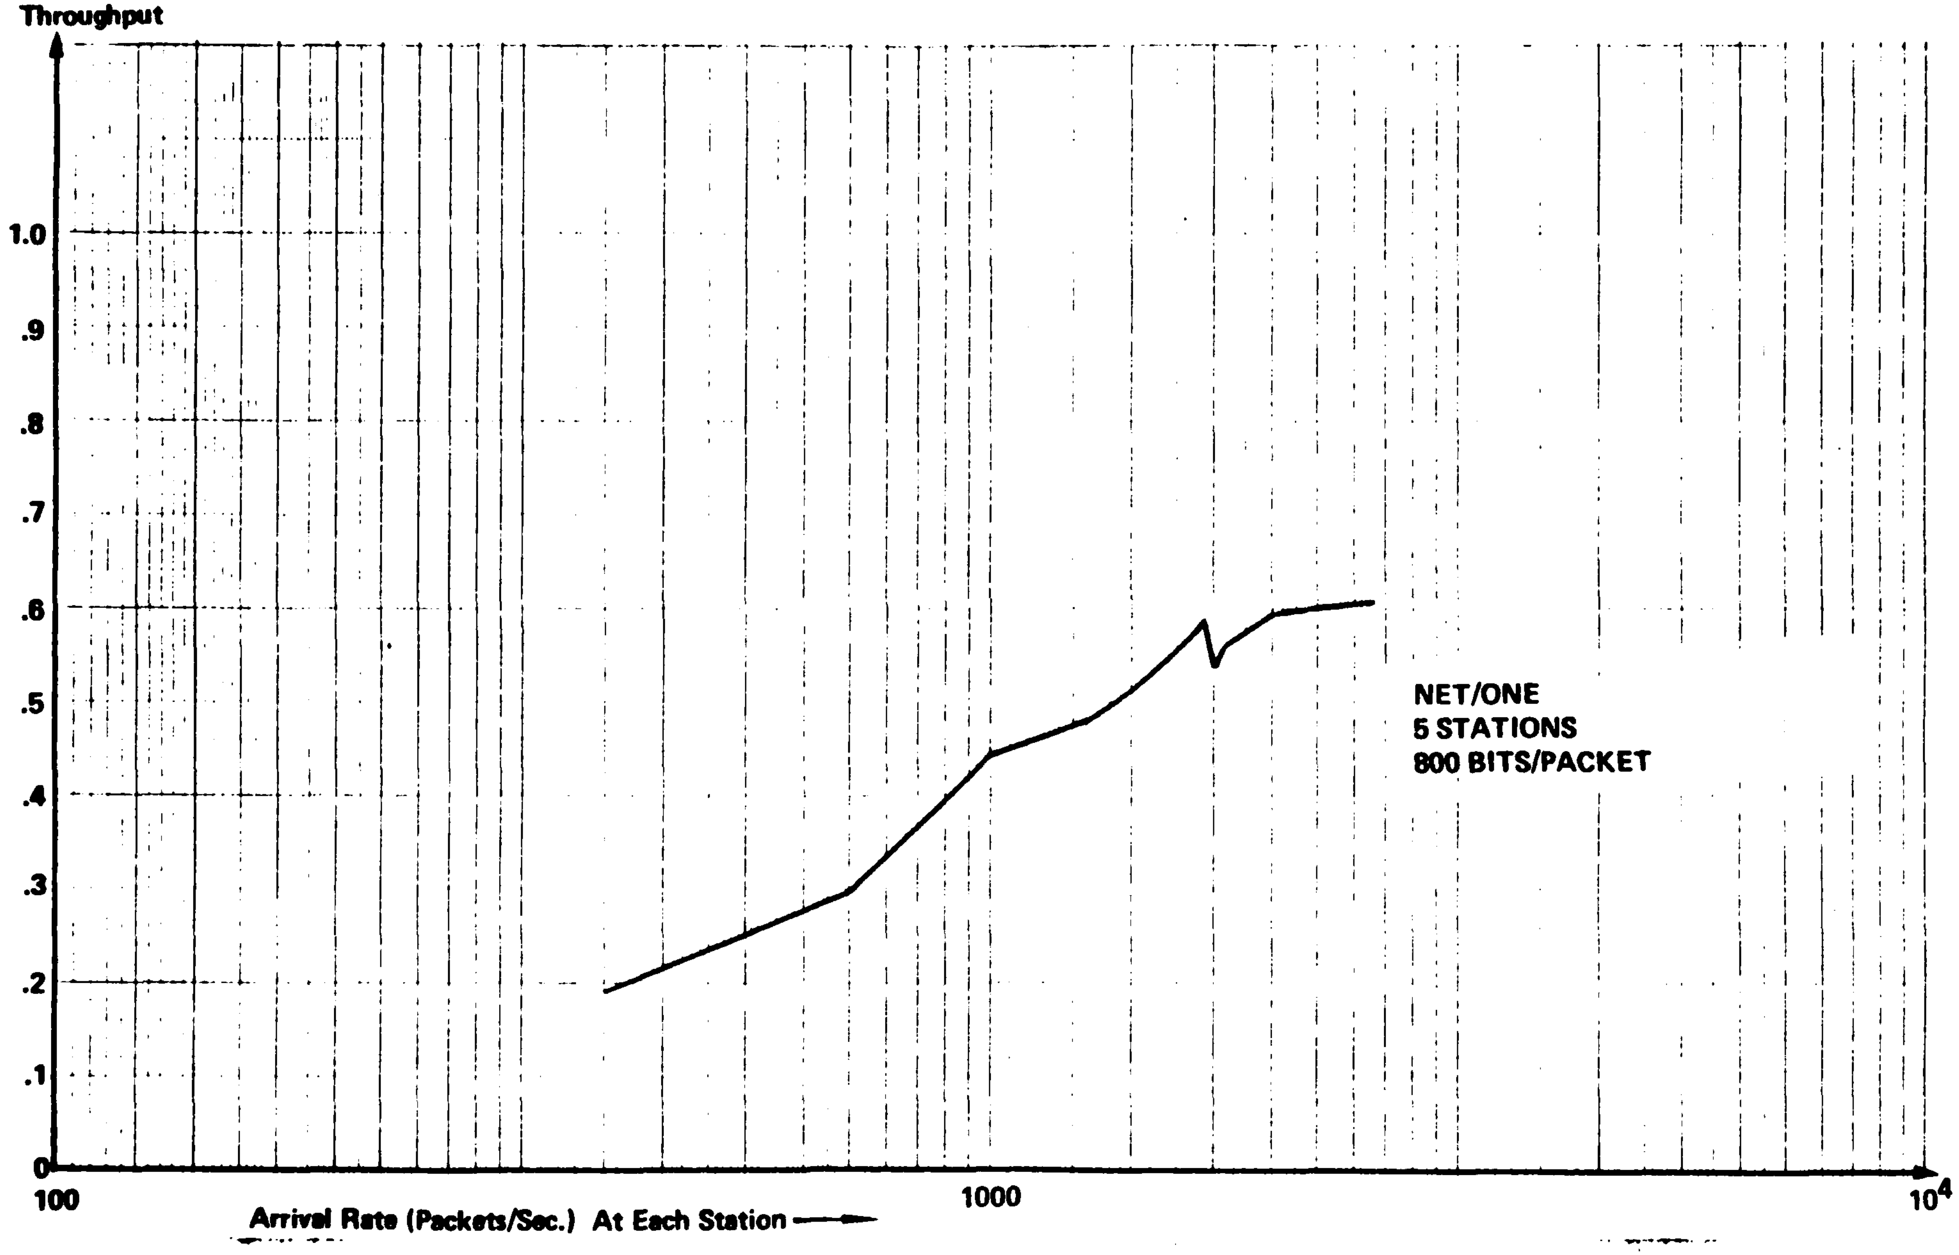

Figure 9. Throughput Versus Packet Arrival Rate for "Ethernet". Levels of Net/One.

```
                            Percentage of total packets
                  10   20   30   40   50   60   70   80   90   100
           1  *
           2  ***********************
           3  ***************************
           4  *
           5
           6
           7
           8
           9
          10
          11
          12
          13
          14
          15
          16
          17
          18
          19
          20
          21
          22
          23
          24
          25
          26
NIU       27
data      28                          No. of ports : 1
bytes     29                          Char rate : 109 char/sec
per       30                          Board : APB
packet    31                          Duration of expt : 5 min
          32                          I/O buffer size : 64 bytes
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                  10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 10. Packet Size Distribution for One Port on APB Active at Character Rate of 109 Per Second.

```
                           Percentage of total packets
                10   20   30   40   50   60   70   80   90   100
          1  ●●
          2  ●●●●●●●●●●●
          3  ●●●●
          4  ●●●
          5  ●●●●
          6  ●●
          7  ●●
          8  ●●●●●
          9  ●●
         10  ●●●●
         11  ●●●●
         12  ●
         13  ●●●
         14  ●
         15  ●
         16  ●●
         17  ●
         18  ●
         19  ●
         20  ●
         21  ●
         22  ●
         23  ●
         24  ●
         25  ●
         26  ●
NIU      27  ●
data     28  ●                          No. of ports : 4
bytes    29  ●                          Char rate : 109 char/sec
per      30  ●                          Board : APB
packet   31  ●                          Duration of expt : 5 min
         32  ●                          I/O buffer size : 64 bytes
         33  ●
         34  ●
         35  ●
         36  ●
         37  ●
         38  ●
         39  ●
         40  ●
         41  ●
         42  ●
         43  ●
         44  ●
         45  ●
         46  ●
         47  ●
         48  ●
         49  ●
         50  ●
         51  ●
         52  ●
         53  ●
         54  ●
         55  ●
         56  ●
         57  ●
         58  ●
                10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 11. Packet Size Distribution for Four Ports on APB, All Active at Rate of 109 Characters per Second.

Figure 12

Packet Size Distribution for One Port on NPB Active at Rate of 109 Characters per second

```
                        Percentage of total packets
                10   20   30   40   50   60   70   80   90   100
           1
           2 ●●●●●●●●●●●●●●●●●●●●●●●●●●
           3 ●●●●●●●●●●●●●●●●●●●●●●●●●
           4
           5
           6
           7
           8
           9
          10
          11
          12
          13
          14
          15
          16
          17
          18
          19
          20
          21
          22
          23
          24
          25
          26
NIU       27
data      28                             No. of ports : 1
bytes     29                             Char rate : 109 char/sec
per       30                             Board : NPB
packet    31                             Duration of expt: 5 min
          32                             I/O buffer size : 64 bytes
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                10   20   30   40   50   60   70   80   90   100
                        Percentage of total packets
```

Figure 12. Packet Size Distribution for One Port on NPB Active at Rate of 109 Characters per Second.

```
                          Percentage of total packets
           .     10    20    30    40    50    60    70    80    90    100
           1  ●
           2  ●
           3  ●
           4  ●●●●●●●
           5  ●●●●●●●●●●●●●●●●●●●●
           6  ●●●●●●●●●●●
           7  ●
           8  ●●
           9  ●
          10  ●●●
          11  ●●●
          12  ●●
          13  ●
          14  ●
          15  ●
          16  ●
          17  ●
          18
          19
          20
          21
          22
          23
          24
          25
          26
NID       27
data      28                             No. of ports : 4
bytes     29                             Char rate  : 109 char/sec
per       30                             Board : NPB
packet    31                             Duration of expt : 5 min
          32                             I/O buffer size : 64 bytes
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                 10    20    30    40    50    60    70    80    90    100
                          Percentage of total packets
```

Figure 13. Packet Size Distribution for Four Ports on NPB. All Active at Rate of 109 Characters per Second.

```
                            Percentage of total packets
                 10   20   30   40   50   60   70   80   90   100
           1
           2  *
           3  *
           4  *
           5
           6
           7
           8  *************************
           9  *************************
          10
          11
          12
          13
          14
          15
          16
          17
          18
          19
          20
          21
          22
          23
          24
          25
          26
#ID       27
data      28                                 No. of ports : 1
bytes     29                                 Char rate : 291 char/sec
per       30                                 Board : APB
packet    31                                 Duration of expt : 5 min
          32                                 I/O buffer size : 64 bytes
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                 10   20   30   40   50   60   70   80   90   100
                            Percentage of total packets
```

Figure 14. Packet Size Distribution for One Port on APB, Active at Rate of 291 Characters per Second.

```
                           Percentage of total packets
                10   20   30   40   50   60   70   80   90   100
          1  ●
          2  ●
          3  ●
          4  ●
          5
          6
          7  ●
          8  ●
          9  ●●●
         10  ●●●
         11  ●
         12  ●
         13  ●
         14
         15  ●
         16
         17
         18
         19
         20
         21
         22
         23
         24  ●
         25
         26
NID      27
data     28  ●                          No. of ports : 4
bytes    29                             Char rate : 291 char/sec
per      30                             Board : APB
packet   31                             Duration of expt : 5 min
         32                             I/O buffer size : 64 bytes
         33  ●
         34
         35  ●
         36  ●
         37
         38
         39
         40
         41  ●
         42  ●
         43  ●
         44  ●
         45
         46
         47
         48
         49  ●
         50  ●
         51  ●
         52
         53
         54  ●
         55
         56
         57  ●
         58  ●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●
                10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 15. Packet Size Distribution for Four Ports on APB, All Active at Rate of 291 Characters per Second.

```
                           Percentage of total packets
                10   20   30   40   50   60   70   80   90   100
           1
           2
           3
           4
           5
           6  *
           7  ************************
           8  *
           9  ************
          10  **************
          11
          12
          13
          14
          15
          16
          17
          18
          19
          20
          21
          22
          23
          24
          25
          26
NIO       27
data      28                           No. of ports : 1
bytes     29                           Char rate : 291 char/sec
per       30                           Board : NPB
packet    31                           Duration of expt : 5 min
          32                           I/O buffer size : 64 bytes
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 16. Packet Size Distribution for One Port on NPB, Active at Rate of 291 Characters per Second.

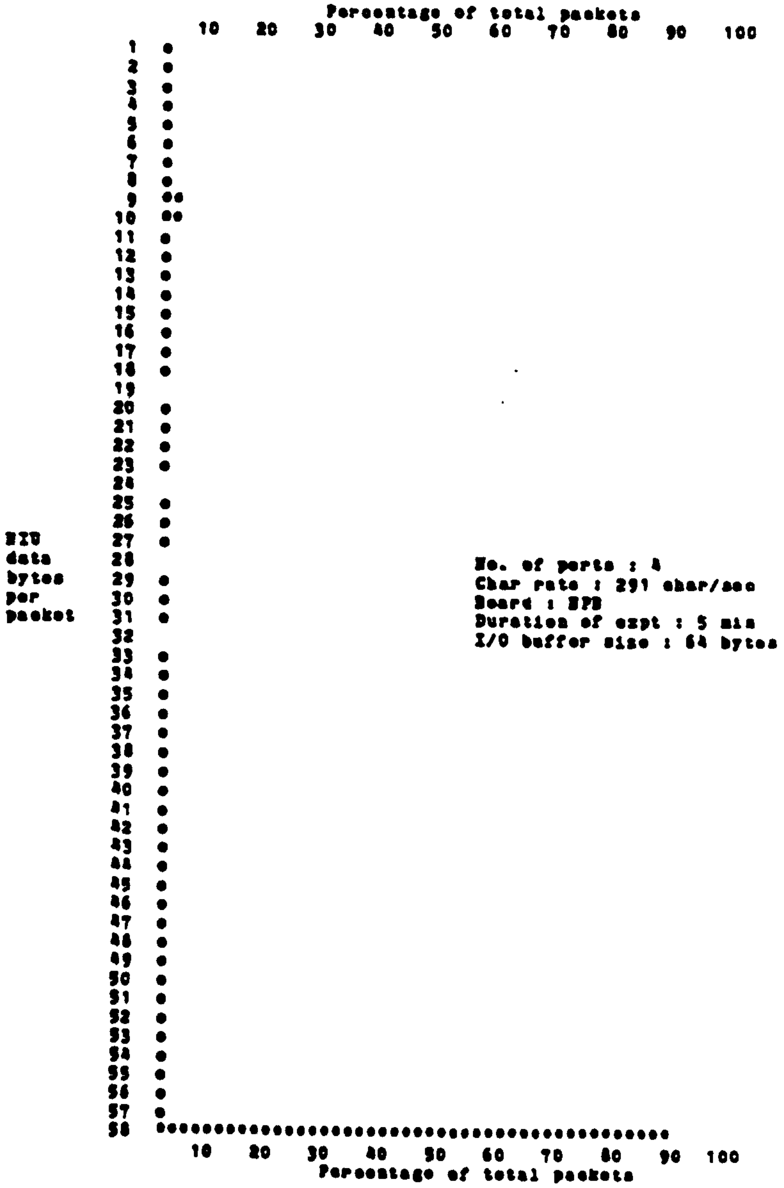

Figure 17. Packet Size Distribution for Four Ports on NPB, All Active at Rate of 291 Characters per second.

```
                              Percentage of total packets
                 10    20    30    40    50    60    70    80    90    100
          1  ●
          2  ●
          3
          4
          5
          6
          7
          8
          9  ●
         10  ●●
         11  ●
         12  ●●●
         13  ●
         14  ●●●●●●●
         15  ●●●●
         16  ●●
         17  ●
         18  ●
         19  ●●●
         20  ●
         21
         22
         23
         24
         25
         26  ●
NIU      27  ●
data     28  ●●●                           No. of ports : 1
bytes    29  ●●●                           Char rate : 436 char/sec
per      30  ●                             Board : NPB
packet   31                                Duration of expt : 5 min
         32                                I/O buffer size : 64 bytes
         33
         34
         35
         36
         37
         38
         39
         40
         41  ●
         42  ●
         43  ●
         44  ●
         45
         46
         47
         48
         49
         50
         51
         52
         53  ●
         54
         55  ●
         56  ●
         57  ●●●●
         58  ●●●●●●●●●●●●●●●●●●
                 10    20    30    40    50    60    70    80    90    100
                              Percentage of total packets
```

Figure 18. Packet Size Distribution for One Port on NPB, Active at Rate of 436 Characters per Second.

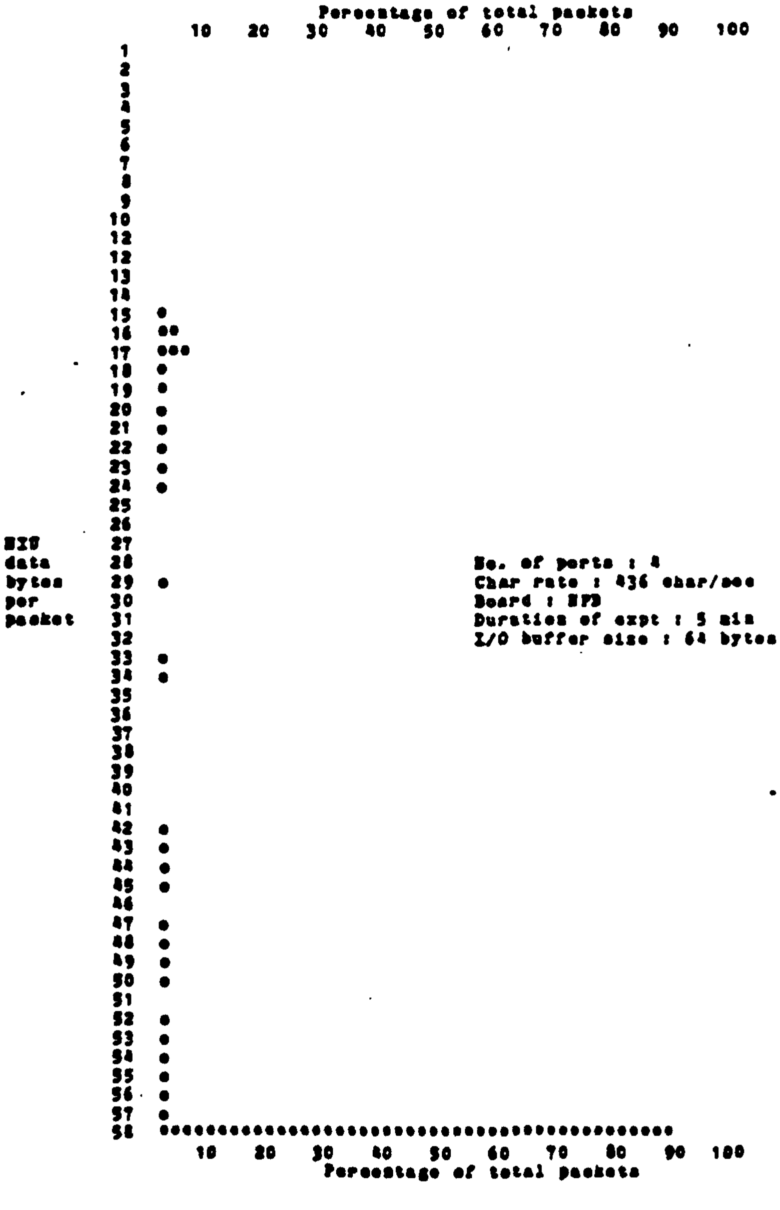

```
                              Percentage of total packets
                 10   20   30   40   50   60   70   80   90   100
              1
              2
              3
              4
              5
              6
              7
              8
              9
             10
             11
             12
             13
             14
             15  *
             16  **
             17  ***
             18  *
             19  *
             20  *
             21  *
             22  *
             23  *
             24  *
             25
             26
SIZE         27
data         28                            No. of ports : 4
bytes        29  *                         Char rate : 436 char/sec
per          30                            Board : NPB
packet       31                            Duration of expt : 5 min
             32                            I/O buffer size : 64 bytes
             33  *
             34  *
             35
             36
             37
             38
             39
             40
             41
             42  *
             43  *
             44  *
             45  *
             46
             47  *
             48  *
             49  *
             50  *
             51
             52  *
             53  *
             54  *
             55  *
             56  *
             57  *
             58  *********************************************
                 10   20   30   40   50   60   70   80   90   100
                              Percentage of total packets
```

Figure 19. Packet Size Distribution for Four Ports on NPB, All Active at Rate of 436 Characters per Second.

```
                           Percentage of total packets
                10   20   30   40   50   60   70   80   90   100
          1  ***********************
          2  **
          3  **
          4  *
          5  **
          6  *
          7  *
          8  *
          9  *
         10  *
         11  *
         12  *
         13  **
         14  *
         15  **
         16  ****
         17  *
         18  **
         19  *
         20  *
         21  *
         22  *
         23  *
         24  *
         25  *
         26  *
NIU      27  *
data     28  *               Packet distribution for a typical
bytes    29  *               working day for a period of 30
per      30  *               minutes
packet   31  *
         32  *
         33  *
         34  *
         35  *
         36  *
         37  *
         38  *
         39  *
         40  *
         41  *
         42  *
         43  *
         44  *
         45  *
         46  *
         47  *
         48  *
         49  *
         50  *
         51  *
         52  *
         53  *
         54  *
         55  *
         56  *
         57  *
         58  ***
         69  *
                10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 20. Packet Size Distribution for a Typical Real Working Load Measured Over a Period of 30 Minutes: Measurement 1.

```
                                  Percentage of total packets
                   10    20    30    40    50    60    70    80    90    100
            1  ●●●●●●●●●●●●●●●●●●●●●
            2  ●
            3  ●
            4  ●
            5  ●
            6  ●
            7  ●
            8  ●
            9  ●
           10  ●
           11  ●
           12  ●
           13  ●●●
           14  ●
           15  ●●●
           16  ●●●●●●●
           17  ●
           18  ●●●
           19  ●●
           20  ●
           21  ●
           22  ●
           23  ●
           24  ●
           25  ●
           26  ●
SIZ        27  ●
data       28  ●                 Packet distribution in typical
bytes      29  ●                 working day for a period of
per        30  ●                 30 minutes
packet     31  ●
           32  ●
           33  ●
           34  ●
           35  ●
           36  ●
           37  ●
           38  ●
           39  ●
           40  ●
           41  ●
           42  ●
           43  ●
           44  ●
           45  ●
           46  ●
           47  ●
           48  ●
           49  ●
           50  ●
           51  ●
           52  ●
           53  ●
           54  ●
           55  ●
           56  ●
           57  ●
           58  ●●
           69  ●
                   10    20    30    40    50    60    70    80    90    100
                                  Percentage of total packets
```

Figure 21. Packet Size Distribution for a Typical Real Working Load Measured Over a Peroid of 30 Minutes: Measurement 2.

```
                           Percentage of total packets
                 10   20   30   40   50   60   70   80   90   100
           1  ********************************
           2  *
           3  *
           4  *
           5  *
           6  *
           7  *
           8  *
           9  *
          10  *
          11  *
          12  *
          13  ***
          14  *
          15  *
          16  **
          17  *
          18  ***
          19  **
          20  *
          21  *
          22
          23
          24
          25
          26
          27
NIU       28            Typing rate with no load at the
data      29            background
bytes     30
per       31
packet    32
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                 10   20   30   40   50   60   70   80   90   100
                           Percentage of total packets
```

Figure 22. Packet Size Distribution for a Typical Typist with No Background Load.

```
                              Percentage of total packets
                  10   20   30   40   50   60   70   80   90  100
           1   *******************************************
           2   **
           3   *
           4   *
           5   *
           6   *
           7   *
           8   *
           9   *
          10   *
          11   *
          12   *
          13   *
          14   *
          15   *
          16   *
          17   *
          18   *
          19   *
          20   *
          21   *
          22   *
          23
          24
          25
          26
NIU       27
data      28                Typing rate with 40% load of raw
bytes     29                packets of sizes 128 bytes (1024 bits)
per       30                at the background
packet    31
          32
          33
          34
          35
          36
          37
          38
          39
          40
          41
          42
          43
          44
          45
          46
          47
          48
          49
          50
          51
          52
          53
          54
          55
          56
          57
          58
                  10   20   30   40   50   60   70   80   90  100
                              Percentage of total packets
```

Figure 23. Packet Size Distribution for a Typical Typist with a 40% Load of 128 Byte Packets also Applied.

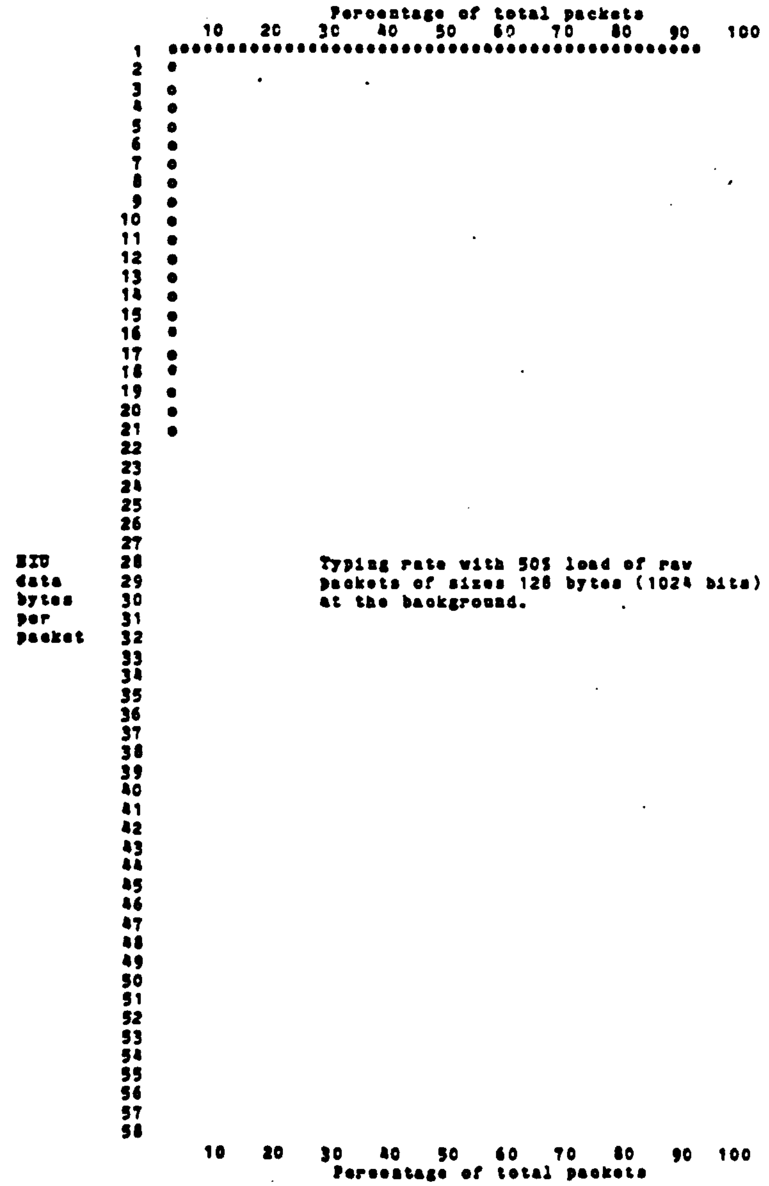

Typing rate with 50% load of raw packets of sizes 128 bytes (1024 bits) at the background.

Figure 24. Packet Size Distribution for a Typical Typist with a 50% Load of 12 128 Byte Packets also Applied.

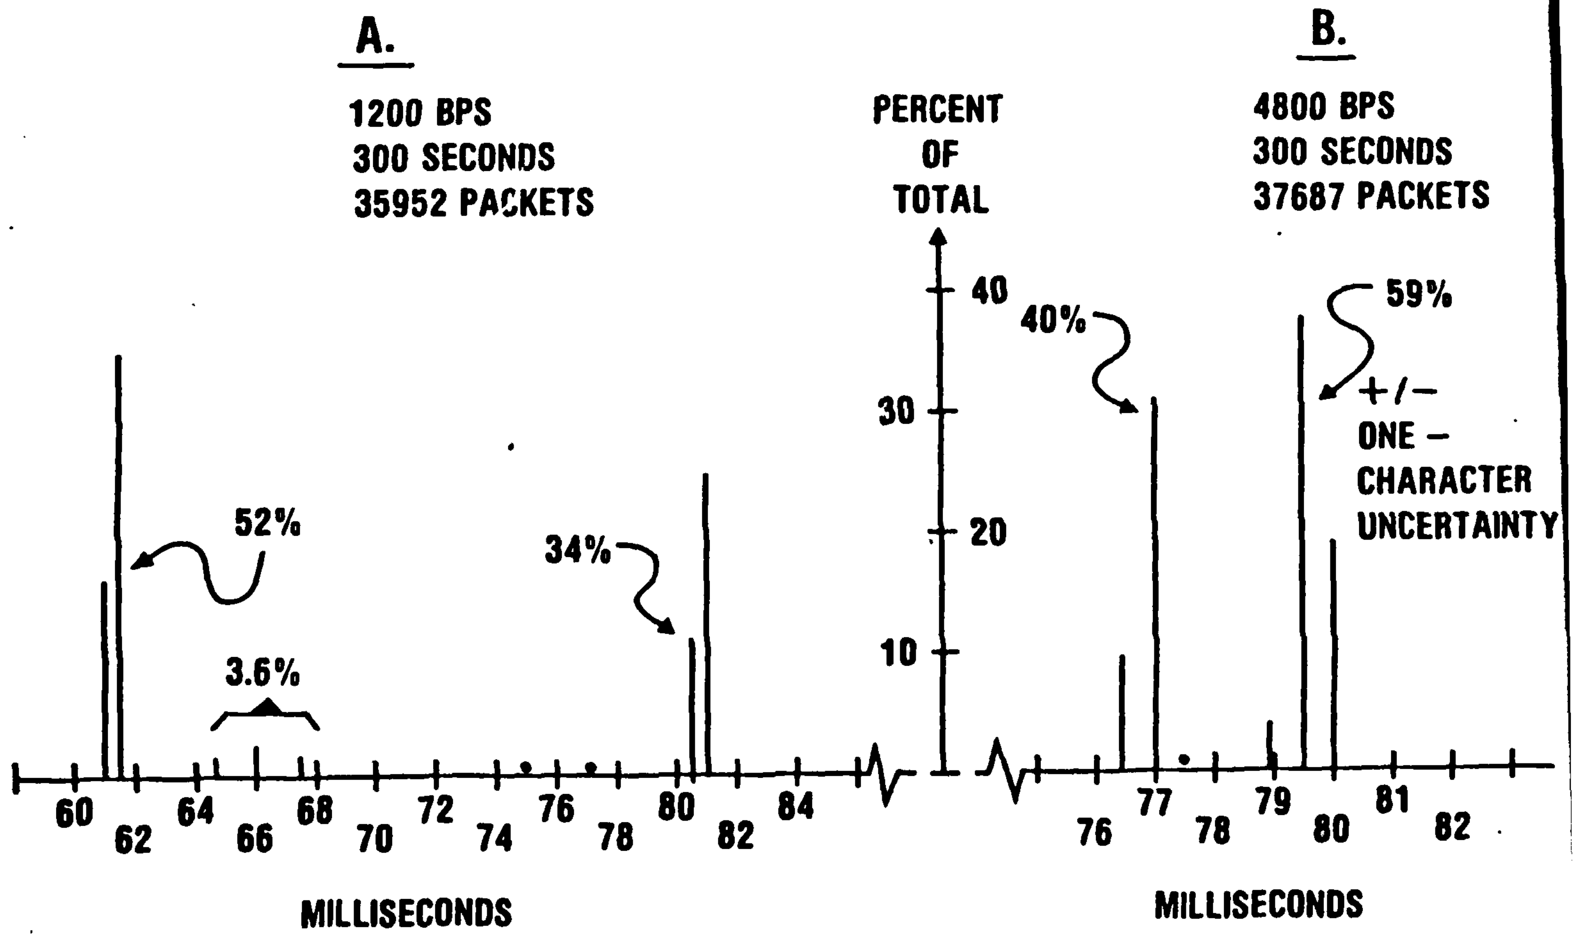

Figure 25. Character Delivery Delay Distributions.

7. EMULATION OF NET/ONE

As noted above, the emulation of Net/One is undertaken in two parts, first the media access and physical layers are emulated and then the higher layers which provide for virtual circuit operation.

7.1 Basic Media Access Operation

As discussed in [Murr83] and[O'RE83], the MLCN facility is structured to emulate the physical and media access layers of a CSMA/CD type local area network in a fairly direct fashion. Provision is made for two primary stations and for a specified number of background stations. Time delays, corresponding to propagation delays over varying lengths of cable, can be programmed between the two primary stations, and between each of the primary stations and the effective location of the background. Average arrival rate, average packet length and number of stations are parameters of the background. The arrival times of packets for the background are structured to be those of a Poisson process. Traffic to the primary stations must be supplied externally.

As discussed in Sections 5 and 6, the basic media access experiment was carried out on the Net/One with a total of five stations each operating to supply 1/5 of the total load. Throughput versus load was measured for the whole network with constant length packets. The case for which a direct emulation can be made used an exponentially distributed time between packet arrivals. The total cable length used by the Net/One was very short, corresponding to 0.212 microsec of delay.

In preparing for the emulation, the average packet length for constant length packets and the average time between packets are specified directly from the Net/One data.

The decision was made to use one primary station and four stations in the background for the emulator. A program was written for the Nova computer to provide constant length packets with an exponentially distributed time between arrivals to the single primary station. The average arrival rate was varied through the same sequence of values used with Net/One.

Either of the two configurations shown in Figure 26 could be used. However, the very small propagation delay (due to the short cable) makes it immaterial which is used.

A study of the Ethernet specifications for the physical layer indictes that there can be several equipment delays of the same order of magnitude as the cable delay for a very short cable. Since instrumentation is not available for meauuring these delays, a discrete-event simulation study was carried out comparing throughput versus arrival rate curves for different values of total delay to those obtained from Net/One. A total delay of 2 msec gives the best results and this value of delay was used for the emulator.

Throughput vs. offered traffic (arrival rate) curves are given in Figure 27 for the emulator in comparison to the curve for Net/One. A good agreement up to large values of load can be noted.

7.2 Virtual Circuit Operation

Emulation of all layers of the Net/One protocol necessary for virtual circuit operation was dicussed briefly in Section 3 and a preliminary configuration for the emulator is given in Figure 3. This approach, planned in the early part of the project, envisioned proramming the essential parts of the higher level Net/One protocols on the Nova computers used at the primary nodes of the emulator.

Unfortunately, subsequent data from the Net/One showed that this approach cannot be used. As is discussed in Section 4, the Net/One does not use straight forward higher layer protocols programmed in software. Instead, as discussed in Section 4, the packetizing process is carried out with a very complicated combination of hardware and software which would be difficult if not impossible to represent with software.

The above facts became clear late in the project and, although several alternatives to the structure for the emulator shown in Figure 3 were considered, no feasible general approach was discovered.

An alternative, and admittedly special purpose, approach to emulating at least a part of the Net/One behavior was chosen. Characteristics of Net/One were measured, as discussed in Sections 5 and 6, under conditions such that the delay caused by the higher level protocols is essentially constant. Under these conditions, variation of average access delay with load is primarily due to the delays encountered in the physical and media access layers. Under conditions of light load, these delays are negligible in comparison to those of the higher level operations. Unfortunately, heavy loads cannot be achieved with the number of stations we had available on the Net/One.

Two sets of emulation data were taken. The first emulates the packet length and arrival rates for a typical experiment on Net/One and adds background stations with essentially the same packet length and arrival statistics. Figure 28 shows the throughput of one primary station versus the number of background stations under the conditions. Note that adding up to 275 stations does not affect throughput.

A second experiment was designed to increase the background load until a noticable decrease in the throughput of the primary station occured. To accomplish this the average arrival rate to the background stations was increased to 180 packets/second and the number of stations was varied from 20 to 275. The results are shown in Figure 29.

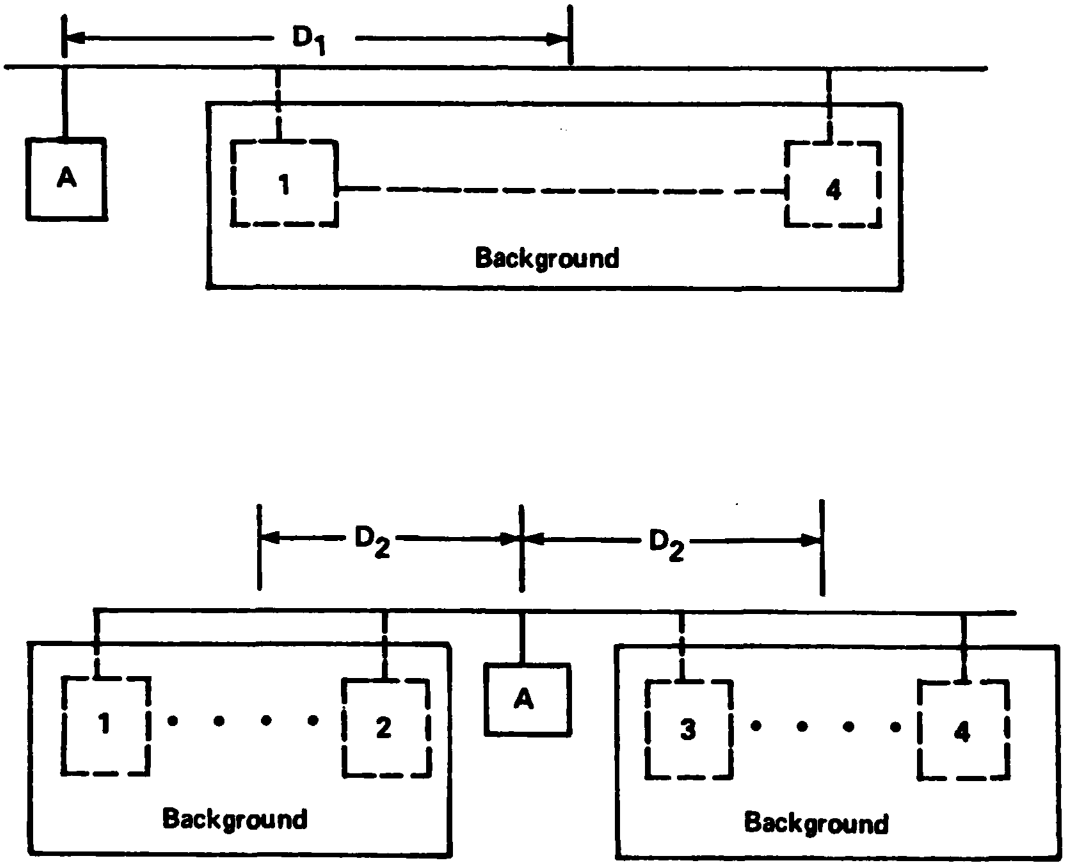

Figure 26. Possible MLCN Emulation Configurations for One Primary Station and Four Background Stations.

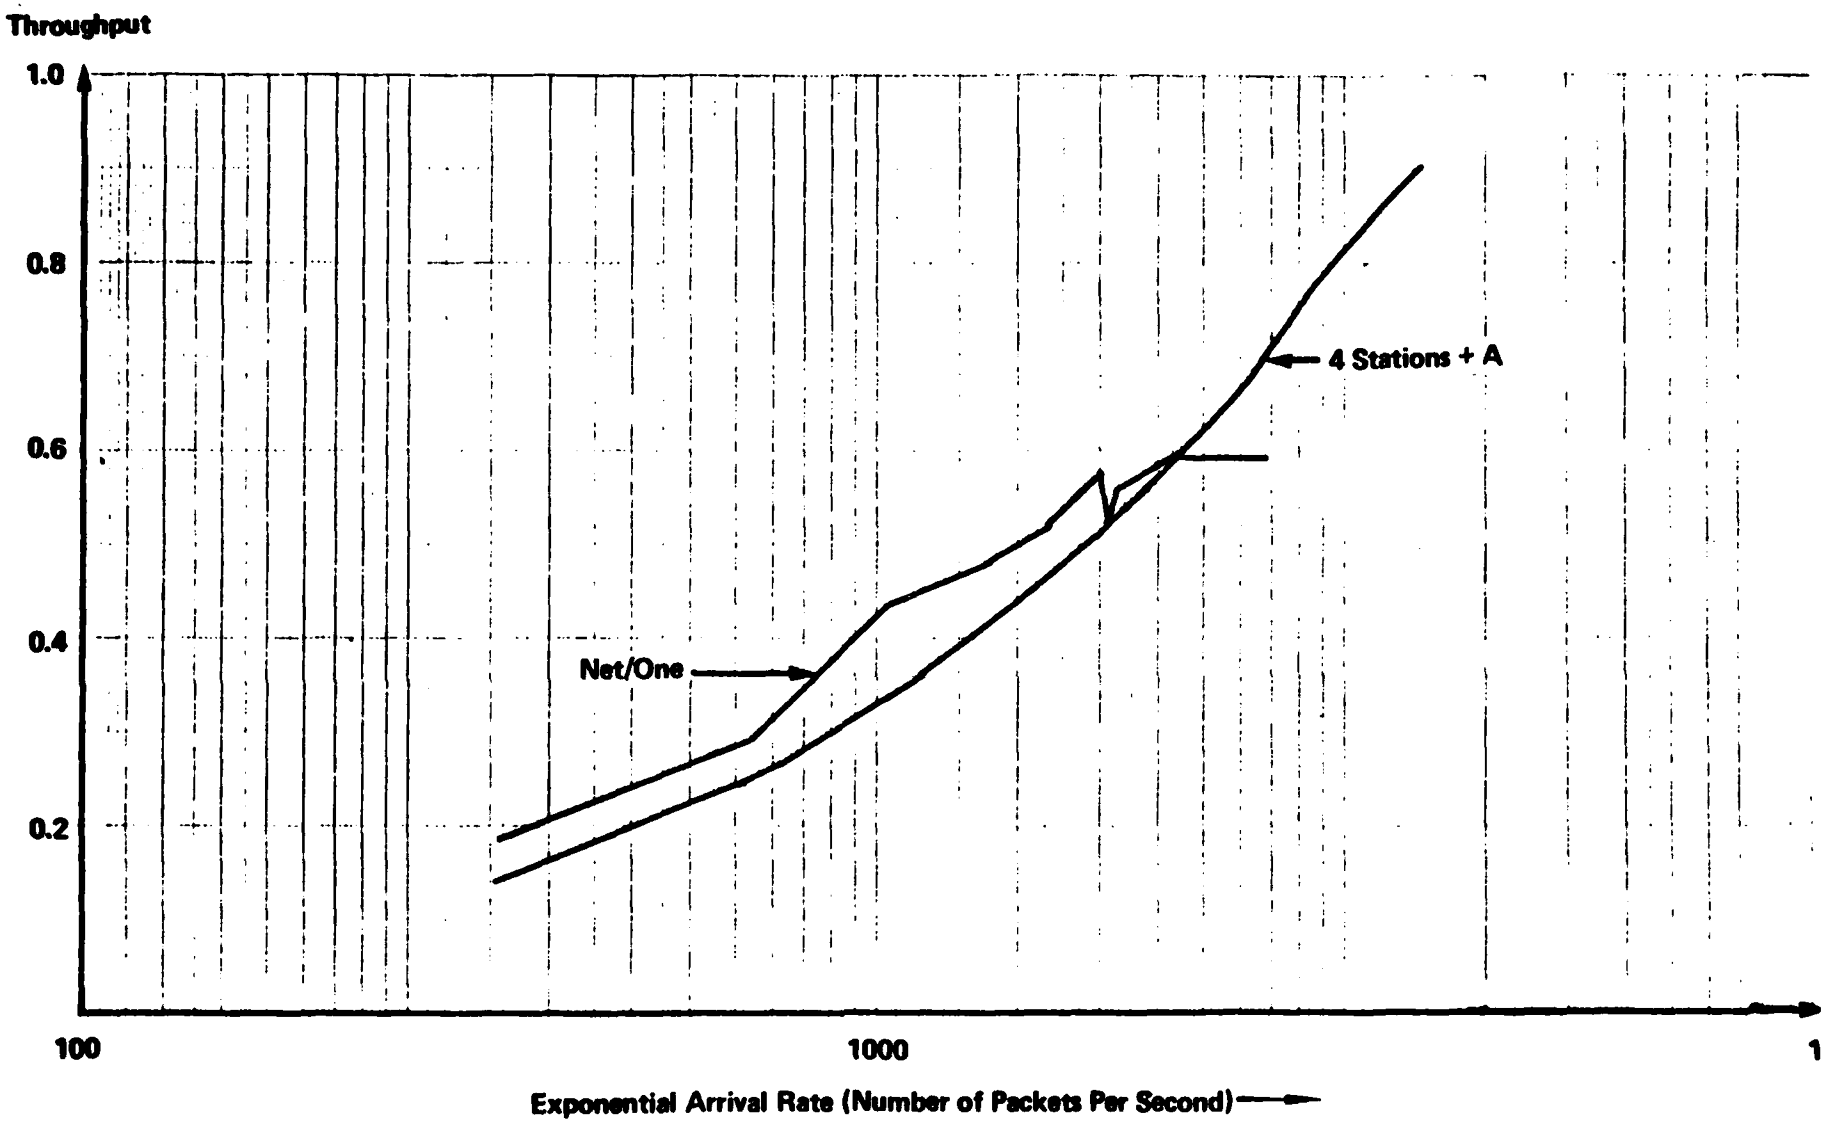

800 Bits Per Packet
Constant Offset of 193 Microsecond

Figure 27. Comparison of Throughput Versus Offered Load for Physical and Media Access Layers of Net/One and Emulation of Net/One.

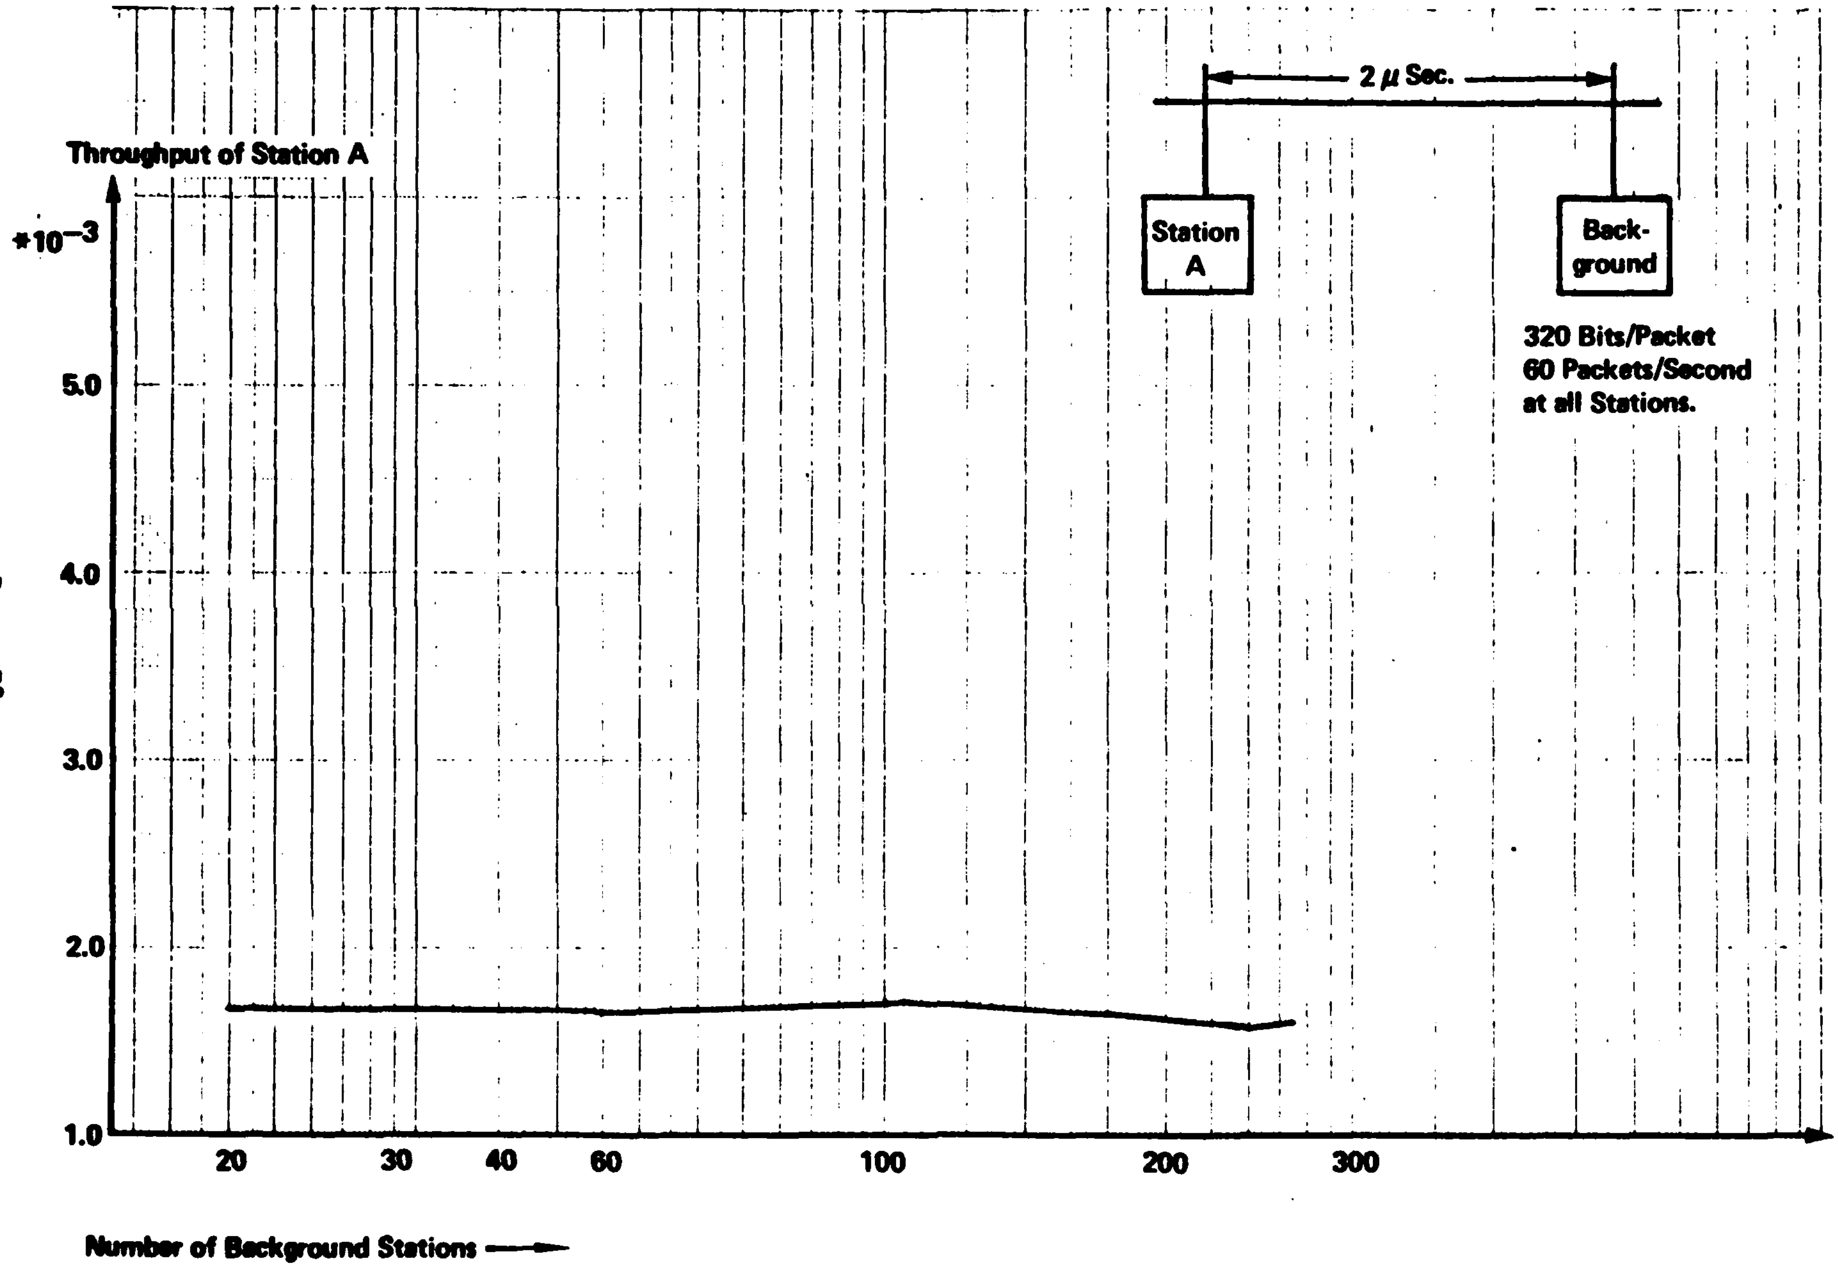

Figure 28. Throughput of One Primary Station Versus Number of Background Stations for Typical Net/One Loading Conditions . Arrival Rate Is 60 Packets Per Second.

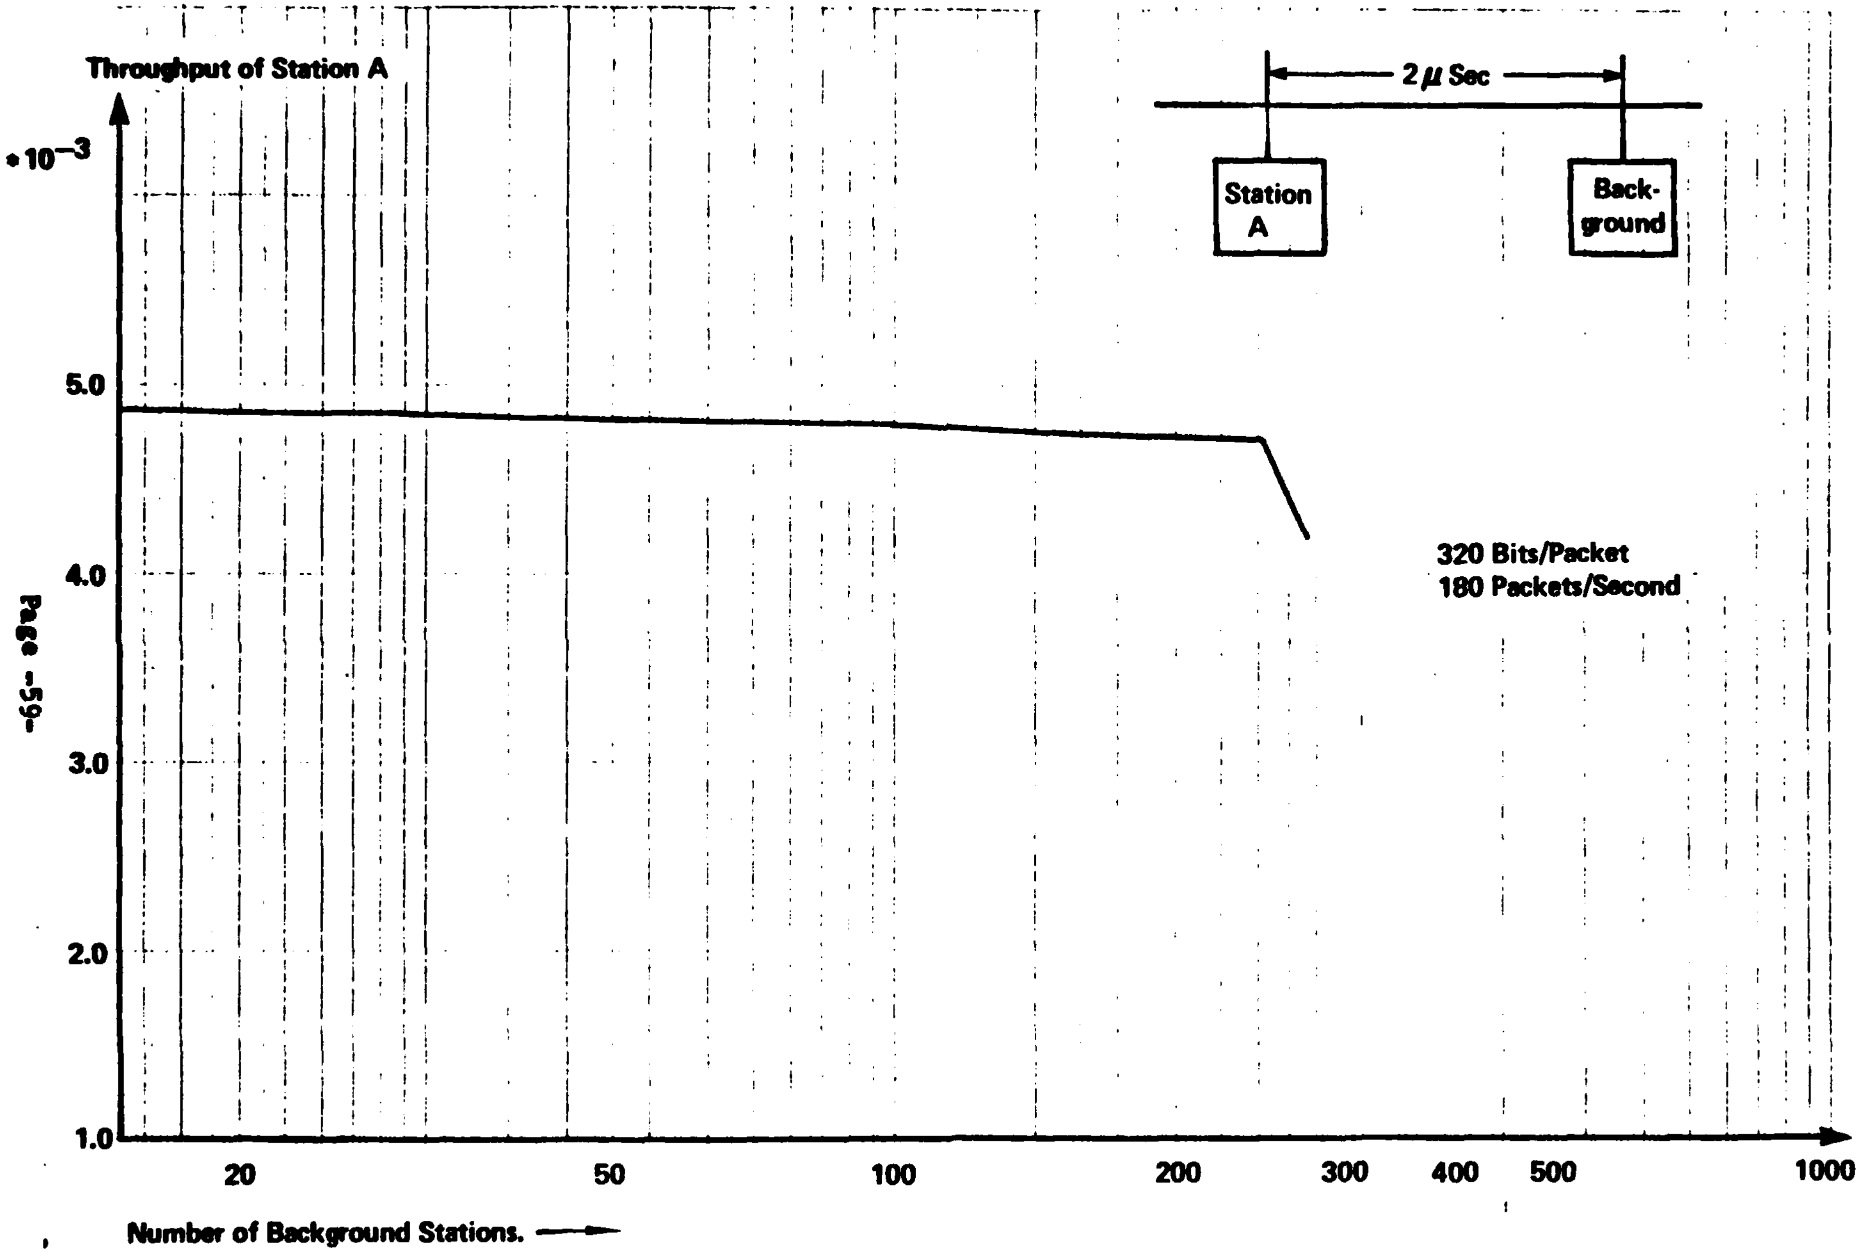

Figure 29. Throughput of Station A Versus Number of Background Stations.

8. Conclusions

The results of the project have been useful, and possibly significant, but not exactly in the directions anticipated. Determining the characteristics of Net/One turned out to be a much more difficult and time consuming task than anticipated. The measured characteristics, when obtained, showed an architecture that is not amenable to emulation on the MLCN as presently structured. The most significant accomplishment of the project thus turned out to be the characterization of Net/One. The majority of this section is devoted to the characterization of Net/One after a brief summary of the emulation results.

8.1 Conclusions with Respect to Emulation of Net/One

The MLCN emulator is tailored to emulate the Ethernet levels of a network. Therefore, not surprisingly, the results from the emulator compared favorably with the results obtained by measurement of the physical and media access layers of Net/One as shown in Figure 26.

Since it is virtually impossible to emulate the higher layers of Net/One on the MLCN as presently structured, the study was limited to determining the effect of background loading on the operation of one Net/One station. The results presented in Figure 28 show that under typical conditions the physical and media access layers have negligible effect on average packet delay. Figure 29 shows, however, that it is possible to load the cable heavily enough so that the physical and media access layers affect delay. Under the conditions used to otain the data for Figure 28, approximately 250 stations are required to affect the throughput of the primary station significantly.

8.2 Observations and Conclusions From Measured Data on Net/One

Measurements show that with a single port operating at 1200 bits per second and no other activity on the network, the delivery delay for one character is approximately 80 milliseconds. Of this total delay, 8.3 milliseconds can be attributed to the serial transfer from the device to the port buffer and another 8.3 milliseconds is required for the output transfer, and on the average, there is an uncertainty of one-half character transfer time, or another 4 ms. The other 59 milliseconds are delivery processing and transfer time delays. There was no other traffic on the network, the cable access delay was negligible, the packet transmission time at 10 MBS was only 67.2 microseconds for a minimum size ETHERNET packet, and the propagation delay is negligible. Therefore, the 59 milliseconds must be attributed to the virtual circuit protocol operations in the NIUs. How this time is divided between SEND and RECEIVE functions is not yet known. (We hope to evaluate that later.)

The throughput and delivery delays achieveable with LAN systems of the type discussed here without loading effects depend primarily on the packet building and buffer management processes. The media transfer speeds, access management and other activities are much less significant. We feel that there are three classes of loading that may have significant effects on throughput and delivery delay, and we are now ready to investigate these more fully. Class One is the activity at the input port and the contention for processing cycles within a single Port Interface Controller of interest. The Second Class is the total combined traffic at all of the ports on a Cable Interface Unit, or the multiplexing that is common in many LANs. The Third Class is the total combined traffic on the media that causes collisions, retrys, and delays in emptying the transmit buffer. These are the delays that prevent the CIU from disposing of the packets and receiving the required acknowledgements. We have presented a sample of Class One and we have studied Class Two. The ability to study Three requires not only an initimate understanding of Classes One and Two, and the capability to measure the parameters associated with them, but also the capability to precisely control the media loading. We believe that One and Two are similar in magnitude and effect and our results indicate they are comparable. We have not been able to study Class Three as of yet but we are moving in that direction. We have observed buffer starvation and the invocation of flow control under only Class One. It is obvious that the inability to empty the TRANSMIT-BUFFER will exergerate the buffer starvation problem. We do not know the magnitude of this effect at this time.

A full understanding of systems such as these and the capability to predict their performance depends almost entirely on the ability to develope a complete timing model such as the one shown in Figure 6 and the ability to quantify the various time intervals involved. It should be noted that although "queues" are utilized, a queuing theory analysis is not totally sufficient due to the real life effects of the interference and interaction between the different processes that cannot be handled in classical queueing theory.

For example, consider the Internal Transfer from the SEND-QUEUE to the TRANSMIT-BUFFER as shown in Figure 6. The packet to be copied or transferred is located in the same memory that is used to accumulate the characters from the ports and assemble the packets. When a Port Interface Chip has received a complete character it must be given memory access to transfer that character into the PORT-BUFFER. These memory cycles given to the Ports are then not available for the Internal Transfer to the TRANSMIT-BUFFER. The result is that increases in the Interface Controller load from higher character input rates, cause delays in the movement of complete packets from the Interface Controller reducing the overall transfer rate just when it needs to increase!

Our analysis shows that, with the exception of the transfers to and from the Ports, <u>all</u> time intervals are random variables whose properties are difficult to predict and must be studied during actual operations. We are at a point now to begin these studies. We have identified the particular measurements needed, and we have a better idea of what test tools are needed to measure them.

REFERENCES

[DIX 80] Digital, Intel, Xerox, "The ETHERNET: A Local Area Network, Data Link and Physical Layer Specifications," Version 1.0, Sep 30, 1980, 82 pp.

[Minn83] Minnich, N.M., and Cotton, C.J., "An Evaluation of Two Unibus ETHERNET Controllers," 8th Conference on Local Area Computer Networks, Oct 83, pp.29-36.

[Murr83] Murray, Dale N., "SPANET: System for Performance Analysis of Networks," IEEE Proceedings, Workshop on Performance and Evaluation of Local Area Networks, WPI, Mar 83, pp. 153-165.

[O'Re83] O'Reilly, P.J.P., Hammond, J.L., Schlag, J.H., and Murray, D.N., "Design of an Emulation Facility for Performance Studies of CSMA-Type Local Networks," 7th Conference on Local Computer Networks, Oct 82, pp.49-57.